Modelling the Danish Real-Time Electricity Market

April 30, 2013

ENFOR

Forecasting and Optimization for the Energy Sector

Contents

1 Introduction

This report considers forecasting models for the real-time electricity market in the Western Danish price area of Nord Pool. The models are developed using a data set covering the period from November 1st 2008 and through January 2010. Three different explanatory variables are considered for all models which are: the forecast system load for the area, the forecast wind power production for the area and the forecast area spot price. All three forecasts are made before noon on the day before delivery. The only case where forecast values are not used as an explanatory variable is when spot prices are used on a horizon shorter than 12 hours. In those cases, the actual spot prices are used since they are always known at least 12 hours in advance. In some cases the ratio between the forecast wind power production and forecast load is used. This ratio is termed wind power penetration. For model estimation, two thirds of the data set is taken for parameter estimation (hereafter termed training period) while the last 1/3 of the data is used as a test period. Consequently, parameters are estimated based on data for the period from November 1st 2008 until September 10th 2009 and subsequently tested on the period from September 11th 2009 until January 31st 2010. When reporting performance measures for the training period, the first 5 weeks of the period are regarded as an initialization period and thus discarded from the measure.

Before considering mathematical models Section 2 describes the problem considered in some detail. The data used is described in Section 3 together with some preliminary analyses. Section 4 describes the methods used and Section 5 present the results of the study and in Section 6 some general conclusions are highlighted.

2 Problem Description

Before we go any further, it is probably appropriate to establish precisely what the real-time market is and what we mean by the imbalance sign and the imbalance penalty. Each day at noon bids for purchase and sale of electricity during the upcoming day are collected by Nord Pool. These hourly bids are subsequently aggregated into a supply and demand curve for which the intersection constitute the spot price, $\pi_t^{(S)}$.

During the hour of operation, differences between the energy contracted and the energy actually produced have to be settled at the real-time market. The producers with production surplus sell their extra production at the down-regulation price, $\pi_t^{(\downarrow)}$, and those short of their contracted volume buy their deficits at the up-regulation price, $\pi_t^{(\uparrow)}$. In the Nordic market model, the regulation prices are bounded by the spot prices at one end so that the down-regulation price can never exceed the spot price and the spot price can never exceed the up-regulation price. Furthermore, in order not to penalize producers that help bring

the system towards stability, only one of the two real-time prices can be different from the spot price at any given hour. Put mathematically, the following holds for the relationship between the three prices at all times:

$$\pi_t^{(\uparrow)} \geq \pi_t^{(S)} \quad \forall t \tag{1}$$

$$\pi_t^{(\downarrow)} \leq \pi_t^{(S)} \quad \forall t \tag{2}$$

$$\pi_t^{(\uparrow)} = \pi_t^{(S)} \quad \text{if } \pi_t^{(\downarrow)} < \pi_t^{(S)} \tag{3}$$

$$\pi_t^{(\downarrow)} = \pi_t^{(S)} \quad \text{if } \pi_t^{(\uparrow)} > \pi_t^{(S)} \tag{4}$$

Now we define the up-and down-regulation penalty, $\psi_t^{(\uparrow)}$ and $\psi_t^{(\downarrow)}$ respectively as

$$\psi_t^{(\downarrow)} = \pi_t^{(S)} - \pi_t^{(\downarrow)} \geq 0 \quad \forall t \tag{5}$$

$$\psi_t^{(\uparrow)} = \pi_t^{(\uparrow)} - \pi_t^{(\uparrow)} \geq 0 \quad \forall t \tag{6}$$

and the imbalance sign as

$$I^{(\pi)} = \text{sign}\{\psi_t^{(\uparrow)} - \psi_t^{(\downarrow)}\} = \begin{cases} 1 & \text{if } \pi_t^{(\uparrow)} > \pi_t^{(S)} \\ 0 & \text{if } \pi_t^{(S)} = \pi_t^{(\uparrow)} = \pi_t^{(\downarrow)} \\ -1 & \text{if } \pi_t^{(\downarrow)} < \pi_t^{(S)} \end{cases} . \tag{7}$$

The rationale for modelling the real-time market in terms of the penalty - and not the price - lies in the application(s) for which the forecasts are intended and the hierarchy of the market prices. Originally, the forecasts were to be used for optimal bidding of a price taking wind power producer into the day-ahead market. It can be shown that, when price risk is disregarded, the bid that optimizes the expected hourly revenue of such a producer is a quantile in the probability distribution of the future production [13]. The quantile value is a function of the expected difference between the spot price and the two real-time prices, i.e. the regulation penalties. When this strategy is extended, first to accommodate the price risk of a price taking producer and subsequently to a price maker, the problem is most easily solved by stochastic programming which in turn calls for stochastic scenarios to be generated [7]. Due to the hierarchical structure of the market, scenarios that respect the regulatory framework of the market are best found via predictions of the regulation penalties and not the prices.

The penalty series, $\psi^{(\uparrow/\downarrow)}$, have the feature that each of them is equal to zero 60-70% of the time. This of is of course because the real-time price and the spot price is equal at these times. For the bidding strategy presented in Zugno et al. [13], only the expected value of the real-time penalties is required or

$$\mathbb{E}\left[\psi_{t+k}^{(\uparrow)}|\chi_t\right] = \mathbb{E}\left[\pi_{t+k}^{(\uparrow)} - \pi_{t+k}^{(S)}|\chi_t\right]$$

$$\mathbb{E}\left[\psi_{t+k}^{(\downarrow)}|\chi_t\right] = \mathbb{E}\left[\pi_{t+k}^{(S)} - \pi_{t+k}^{(\downarrow)}|\chi_t\right] .$$

Since there is an obvious regime shift in the penalties, the expected penalty for each regulation direction can be decomposed, using the law of total probability, into

$$\mathbb{E}\left[\psi_t^{(\uparrow)}\right] = \mathbb{P}(\psi_t^{(\uparrow)} > 0) \cdot \mathbb{E}\left[\psi_t^{(\uparrow)}|\psi_t^{(\uparrow)} > 0\right] + \mathbb{P}(\psi_t^{(\uparrow)} = 0) \cdot 0$$

$$\mathbb{E}\left[\psi_t^{(\downarrow)}\right] = \mathbb{P}(\psi_t^{(\downarrow)} > 0) \cdot \mathbb{E}\left[\psi_t^{(\downarrow)}|\psi_t^{(\downarrow)} > 0\right] + \mathbb{P}(\psi_t^{(\downarrow)} = 0) \cdot 0.$$

This in turn allow us to divide the problem into estimation of the probabilities of regulation in each direction and forecasting of the penalties conditioned upon that there is "active" regulation in that direction.

From a modelling perspective, working with the penalties, opposed to the actual prices has some advantages. For one, the penalties have a constant lower boundary at zero which to some extent can mitigate high positive residuals. More importantly though, the penalty forecasts are applicable at all horizons, both day-ahead and intra-day, combined with the current available information about the spot prices.

In light of this, the focus in the following is on the forecasting of the imbalance sign and the penalty separately. However, Section 5.3 evaluate the unconditional expectations and Section 5.4 considers models estimated on the full data set.

3 Data Analysis & Choice of Data Set for the Penalty Models

3.1 Penalties

The regulation penalties, in either direction, have a quite heavy tailed distributions and also have relatively frequent spikes as can be seen in Figures 1 & 2.

Furthermore, judging from Figure 3, the penalties seem to have a clear diurnal cycle.

In order to avoid parameter estimates that solely focus on the residuals from these spikes, some of the most severe spikes have to be removed from the data set. On the other hand, these spikes are a reality of the market and thus it is desirable to have them represented in the estimation set to some extent. Therefore the a scheme was adopted to remove the observations from the series. The scheme can be outlined as follows:

- Adaptively and robustly fit models of the form

$$
\begin{aligned}
\widehat{\psi}_{t+k|t}^{(\uparrow/\downarrow)} = & \phi_{0,t} + \phi_{1,t} \frac{\widehat{W}_{t+k|t}}{\widehat{L}_{t+k|t}} + \phi_{2,t} \widehat{pi}_{t+k|t}^{(S)} \\
& + \sum_{i \in Shr} \alpha_{i,t} \sin\left(\frac{2\pi i \cdot hr(t+k)}{24}\right) \\
& + \sum_{i \in Chr} \beta_{i,t} \cos\left(\frac{2\pi i \cdot hr(t+k)}{24}\right) \\
& + \sum_{i \in Swd} \gamma_{i,t} \sin\left(\frac{2\pi i \cdot wd(t+k)}{7}\right) \\
& + \sum_{i \in Cwd} \delta_{i,t} \cos\left(\frac{2\pi i \cdot wd(t+k)}{7}\right)
\end{aligned}
\tag{8}
$$

  for the penalties in each direction where $\widehat{W}_{t+k|t}$, $\widehat{L}_{t+k|t}$ and $\widehat{pi}_{t+k|t}^{(S)}$ are the forecast wind power production, load and spot price. The model is estimated using the full data set for various pairs of forgetting factor (λ) and cut-off value (τ).

- Calculate the contribution of each individual observation to the RMSE.

- Exclude the observations that contribute more that 0.5% to the total RMSE.

The model in Eq. (8) was chosen as one that could easily include double seasonality, multiple external variables and yet involved a limited number of parameters to vary. The elements of *Shr*, *Chr*, *Swd* and *Cwd* were found as the ones optimizing the AIC criteria (see e.g. [6]) of

the time-invariant version of Eq. (8) which yielded:

	Shr	*Chr*	*Swd*	*Cwd*
Down	$\{1\}$	$\{2\}$	$\sim$	$\{2,3\}$
Up	$\{2,3\}$	$\{1\}$	$\{1,2\}$	$\{2\}$

Figures 4 & 5 show the cumulative squared residuals for the data period for the up-and-down regulation penalties respectively. In both cases there are considerable jumps in the cumulative sum on several occasions, indicating an exceptionally high residual at that time.

This is more apparent in Figures 6 & 7 where the proportional reduction in *RMSE* resulting from excluding the *n* largest residuals (on the x-axis) is shown. The plot show that the 30 largest residuals (out of 3466) for the up-regulation are accountable for almost 60% of the *RMSE* for the full set and the 100 largest ones make up 73.5% of the full *RMSE*. The same numbers for the down-regulation are less severe or 20% and 47% respectively (out of 4646) but are still quite high.

Figures 8 & 9 show the marginal reduction for each of the observations.

The choice of the 0.5% limit is somewhat arbitrary but one important feature of having this limit is that it is the lowest one for which discarded observations are not parameter specific. In Figure 10 the two time series are plotted again with discarded observations marked. The number of excluded observations for the up-and down-regulation penalties is 37 and 18 respectively.

The same plots as in Figures 1 - 3 only for the series cleaned for the most extremes are shown in Figures 11 - 13. There are still quite a lot of spikes present and the implications of these will be discussed further on in this document. Comparing Figures 3 and 13 show that overall the patterns are similar. However, for up-regulation the range of the diurnal variations has changed markedly.
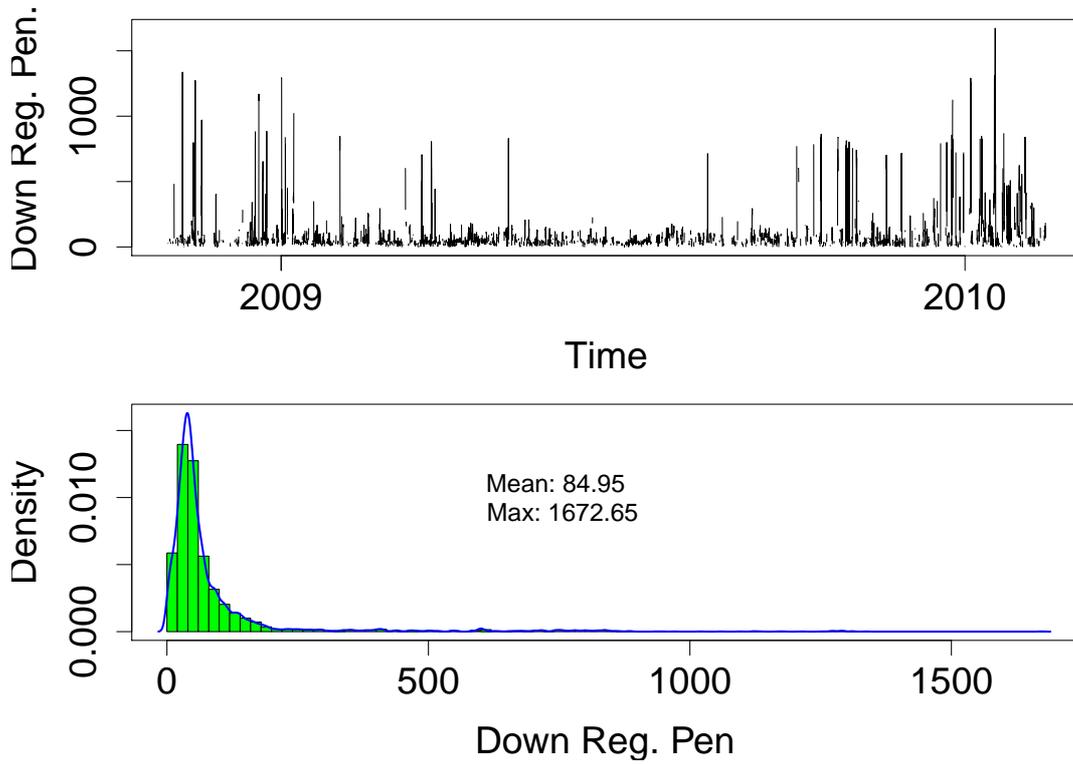
Figure 1: *Time series plot and histogram of the positive down-regulation penalties.*
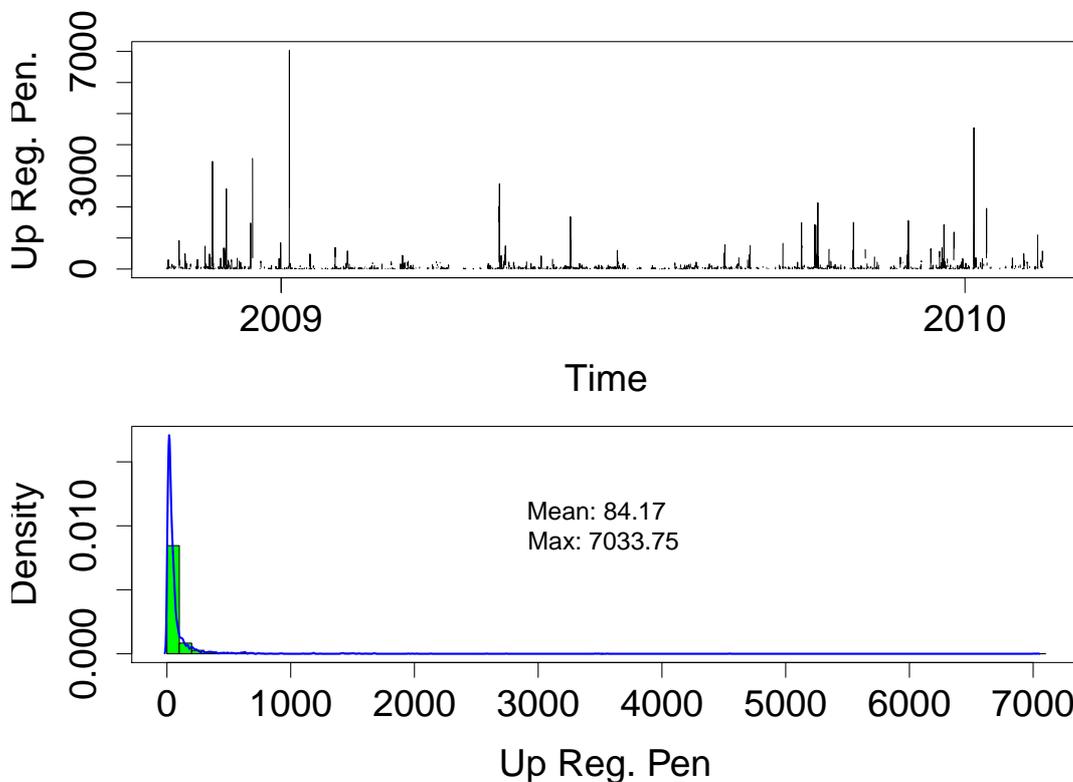

Figure 2: *Time series plot and histogram of the positive up-regulation penalties.*
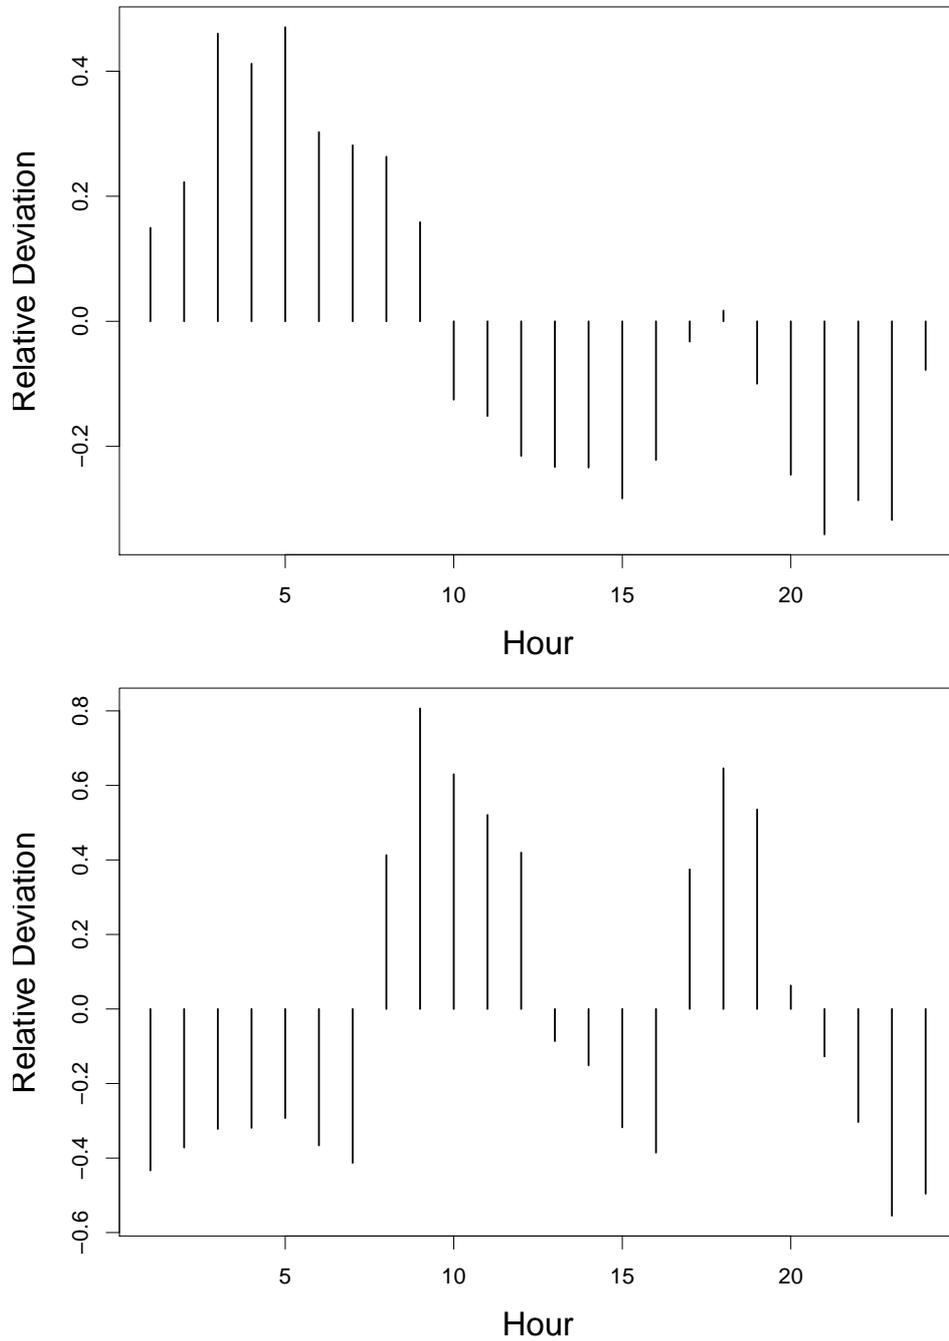
Figure 3: The diurnal variation of the positive down- and up-regulation penalties (top & bottom respectively) given relative to the overall mean. Note that the scale differ between plots.
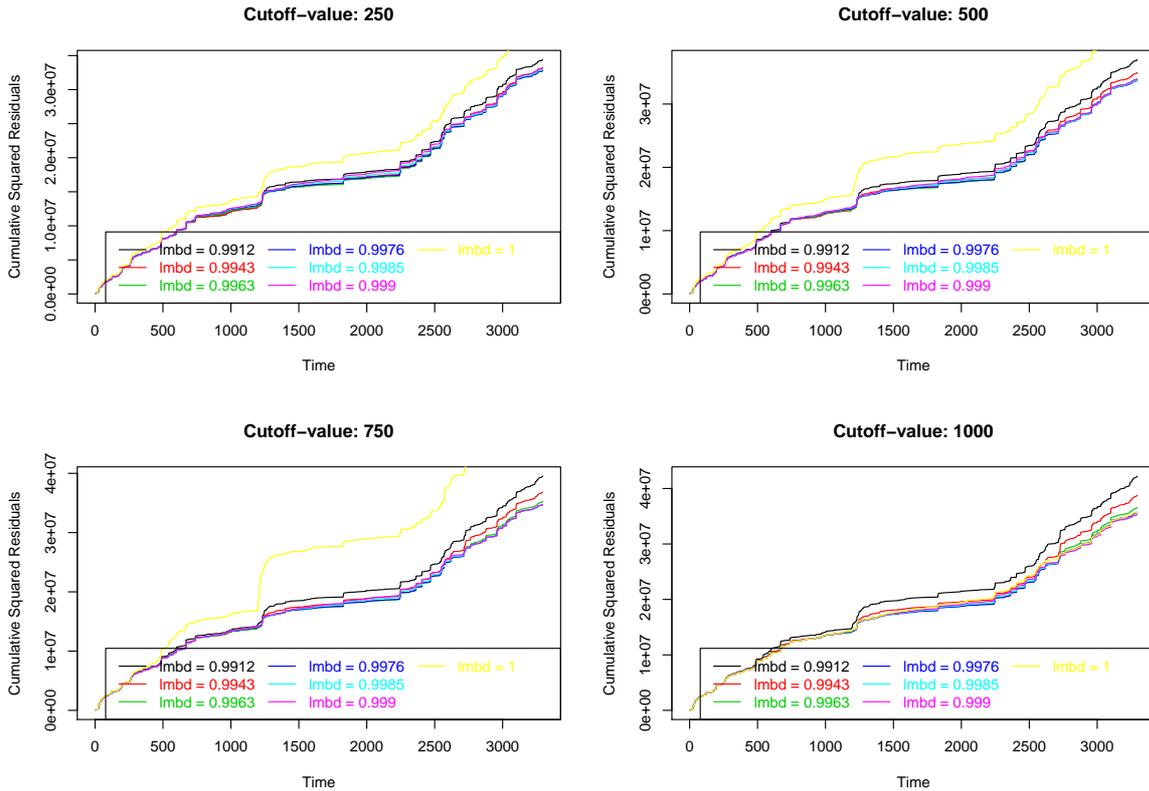
Figure 4: Cumulative squared residual of the up-regulation penalties for different forgetting factors (λ) and cutoff values (τ). Note that the scale differ between plots.

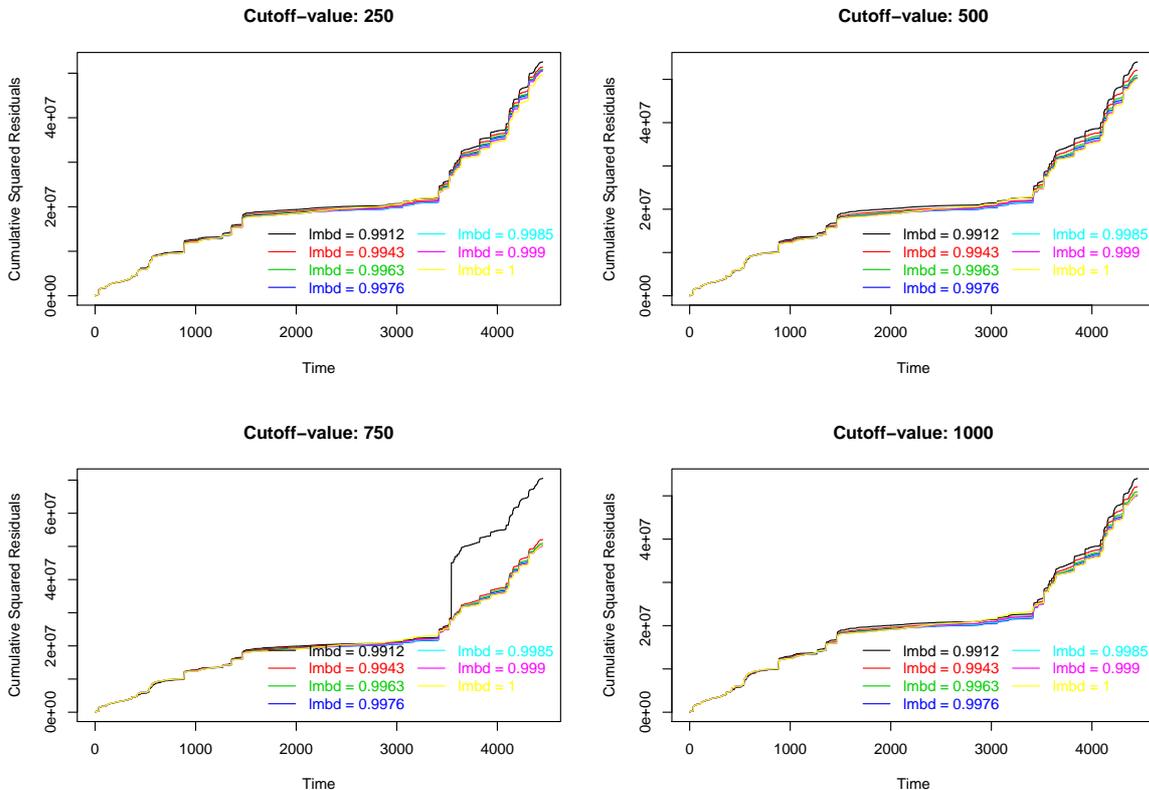*Figure 5: Cumulative squared residual of the down-regulation penalties for different forgetting factors (λ) and cutoff values (τ). Note that the scale differ between plots.*
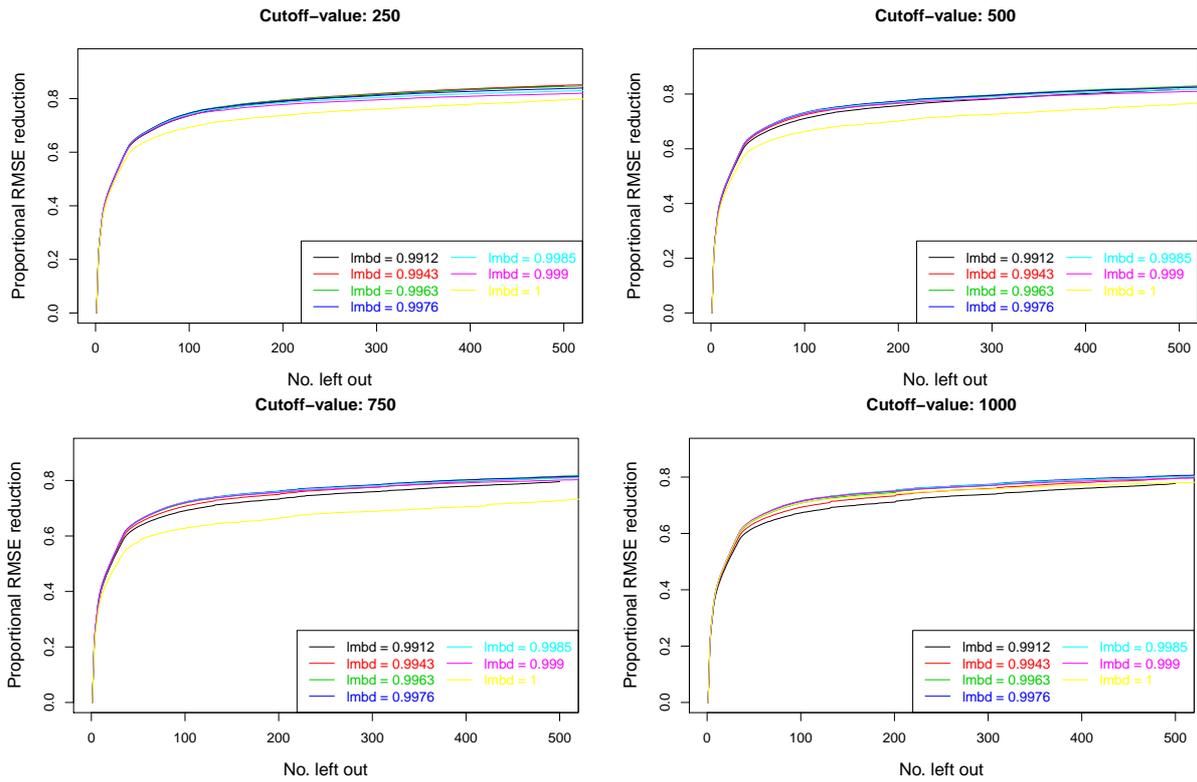
Figure 6: Proportional reduction of the RMSE by removing the n largest residuals for the up-regulation penalties.
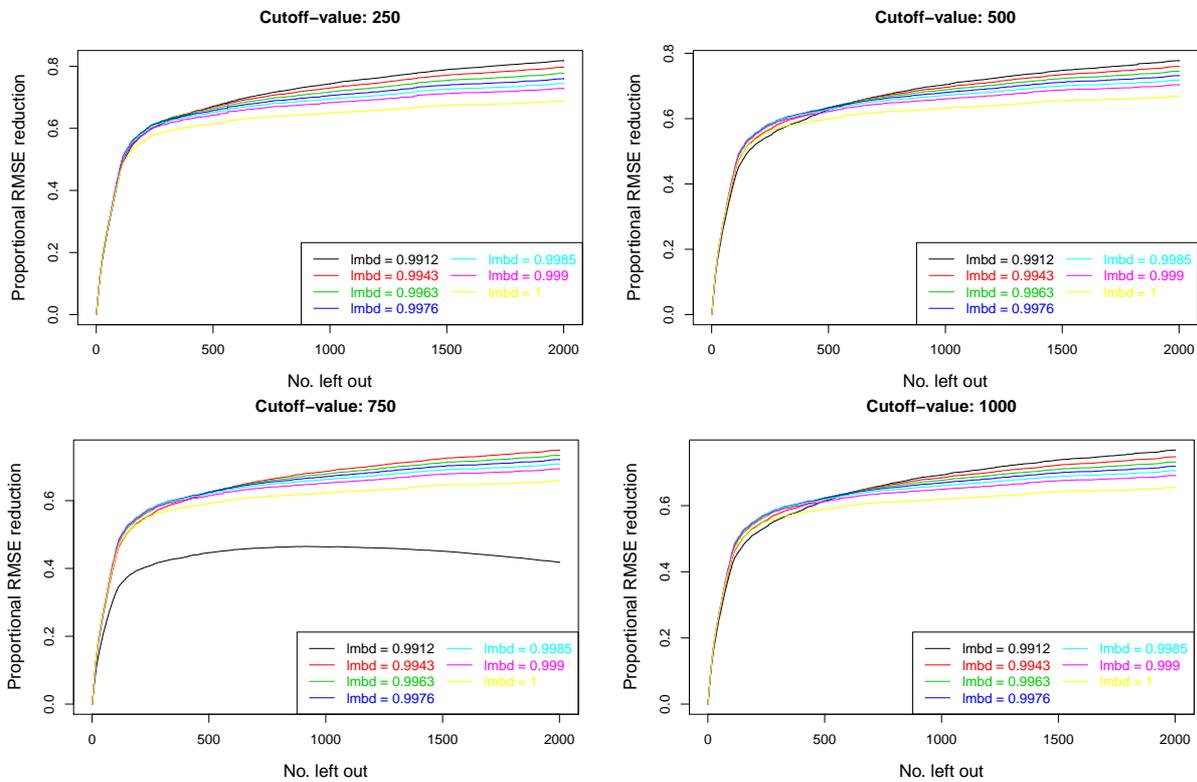

Figure 7: Proportional reduction of the RMSE by removing the n largest residuals for the down-regulation penalties.

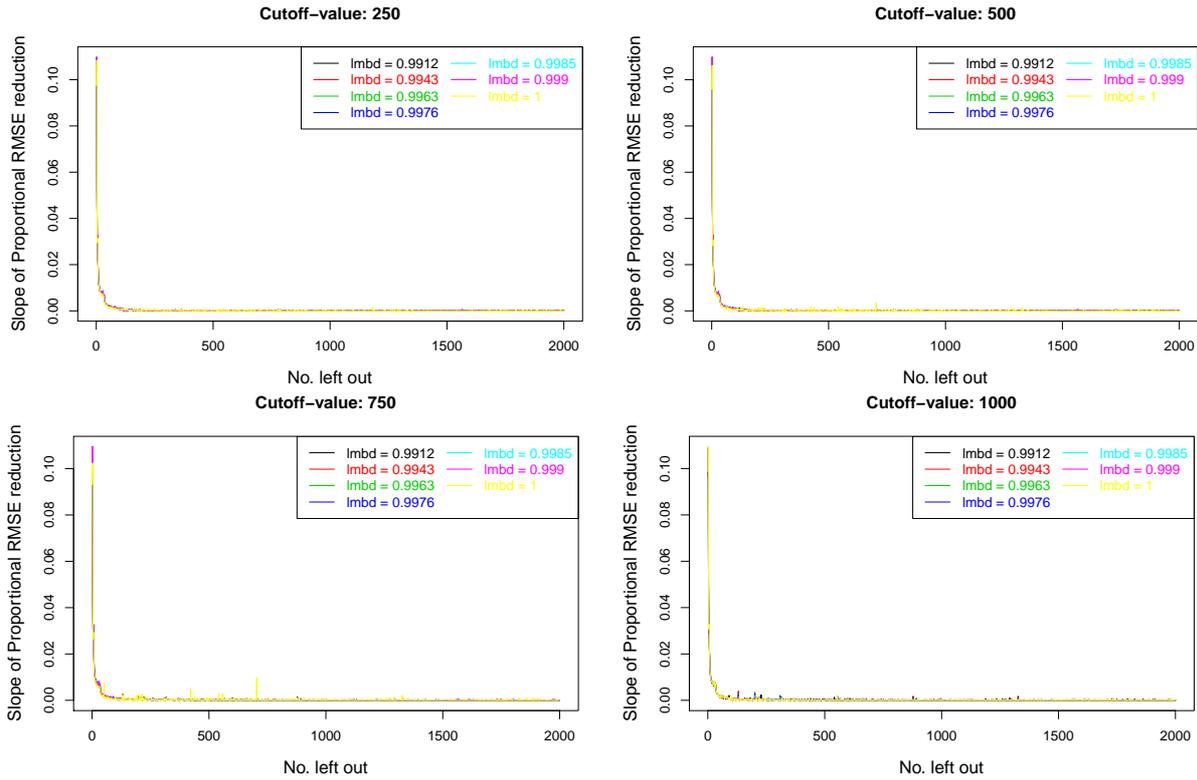*Figure 8: Marginal proportional reduction of the RMSE by removing the nth largest residuals for the up-regulation penalties.*
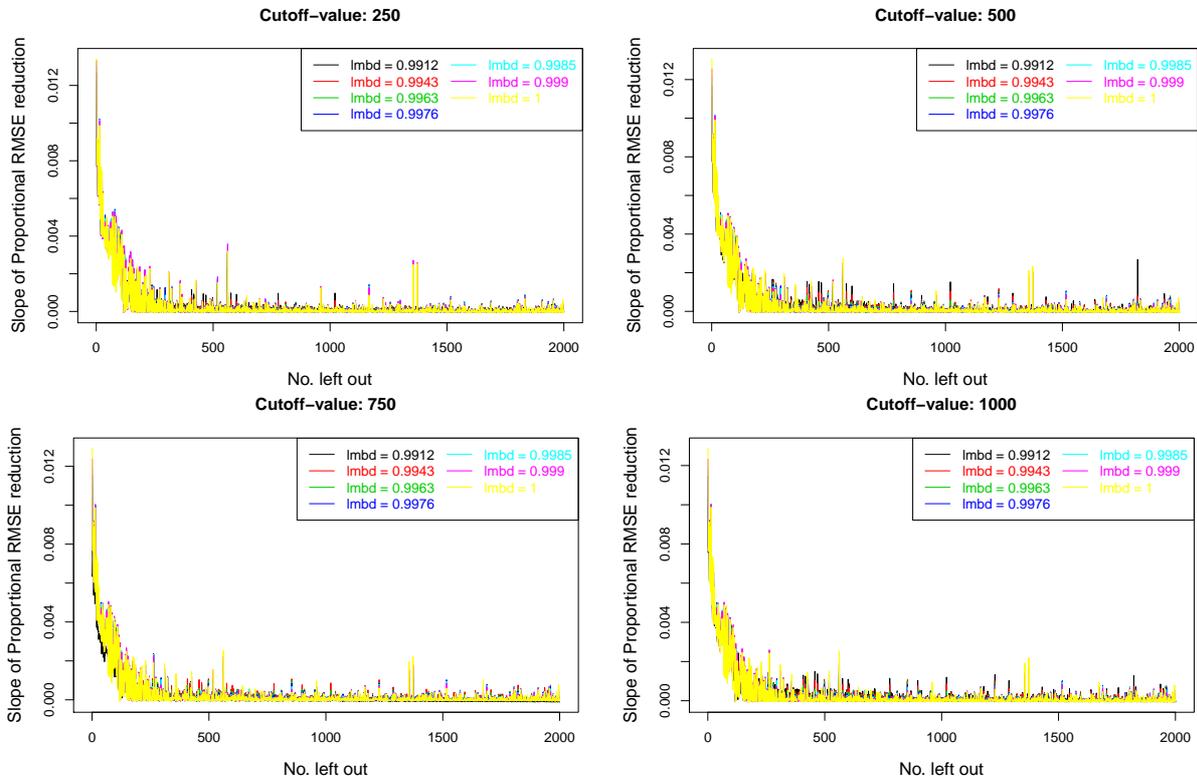
Figure 9: Marginal proportional reduction of the RMSE by removing the nth largest residuals for the down-regulation penalties.
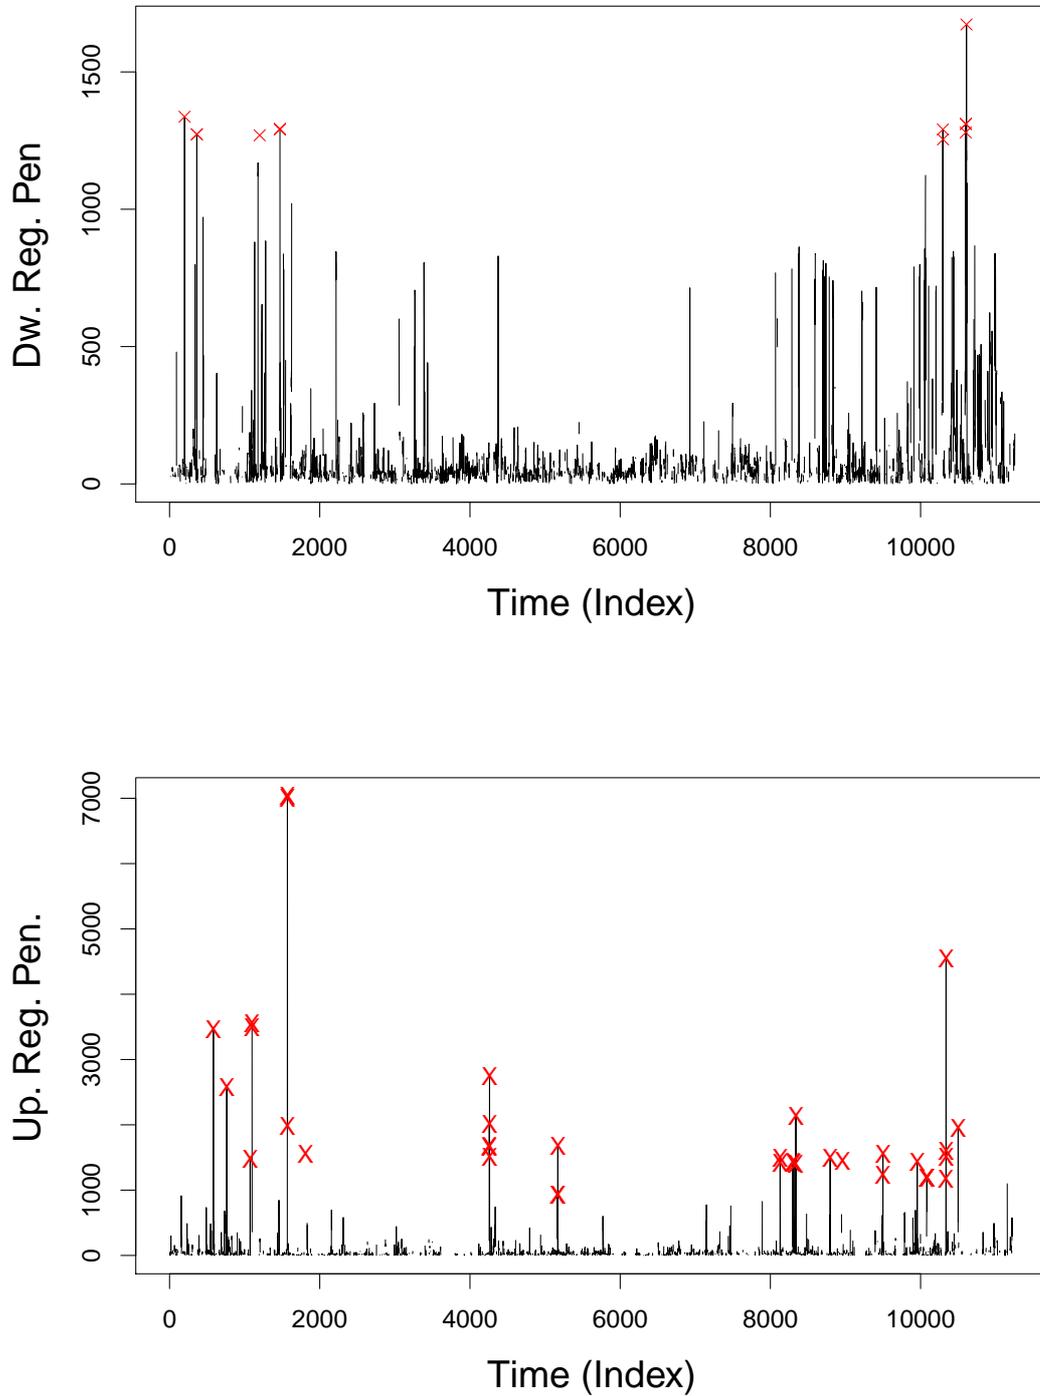
Figure 10: Time series plot of the down-and up-regulation penalties (top and bottom respectively) with the excluded observations marked.

Figure 11: *Time series plot and histogram for the cleaned series of down-regulation penalties.*

Figure 12: *Time series plot and histogram for the cleaned series of up-regulation penalties.*

Figure 13: The diurnal variation of the down-and up-regulation penalties (bottom & bottom respectively) given relative to the overall mean for the cleaned series.

3.2   Imbalance Sign

Due to different conventions for currency exchange rates between Nord Pool and Energinet.dk, price differences less than 0.10 DKK are rounded to zero. Apart from that, the imbalance sign is found by Eq. (7).

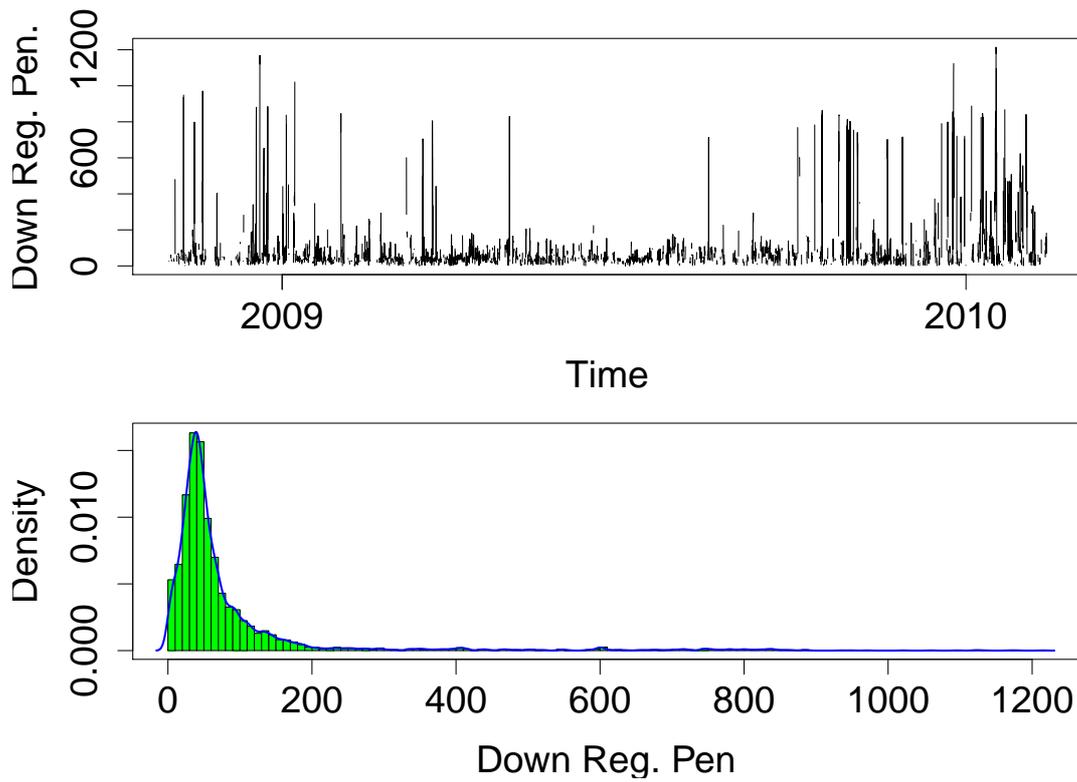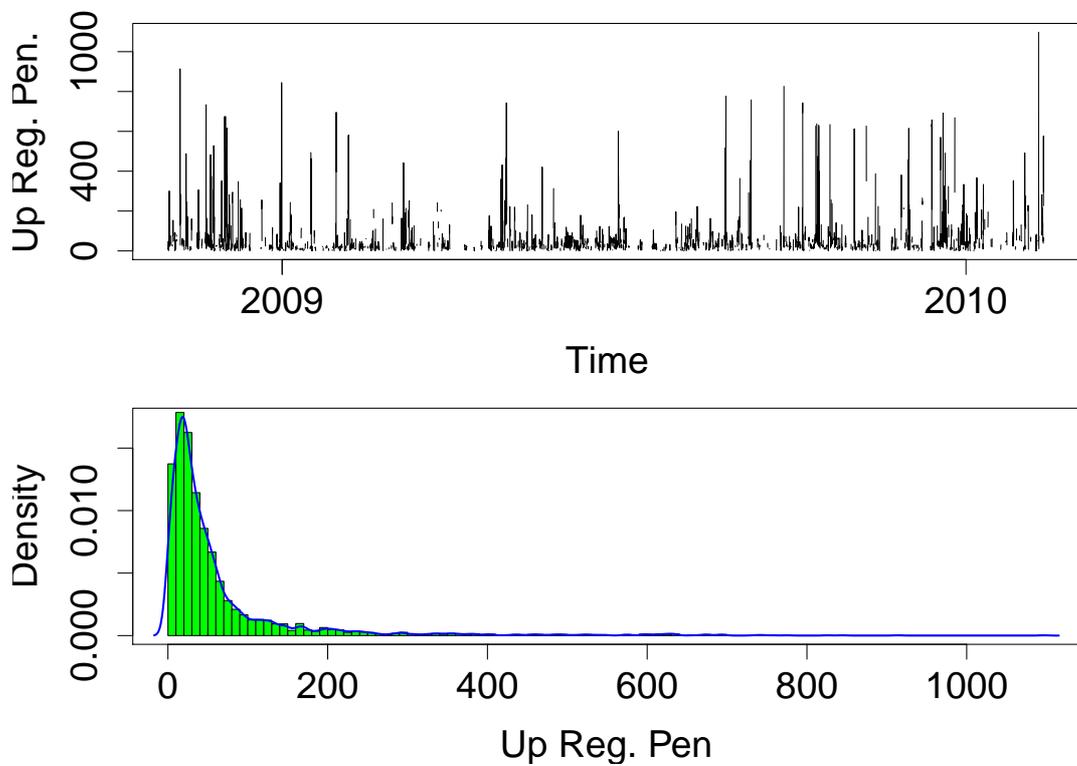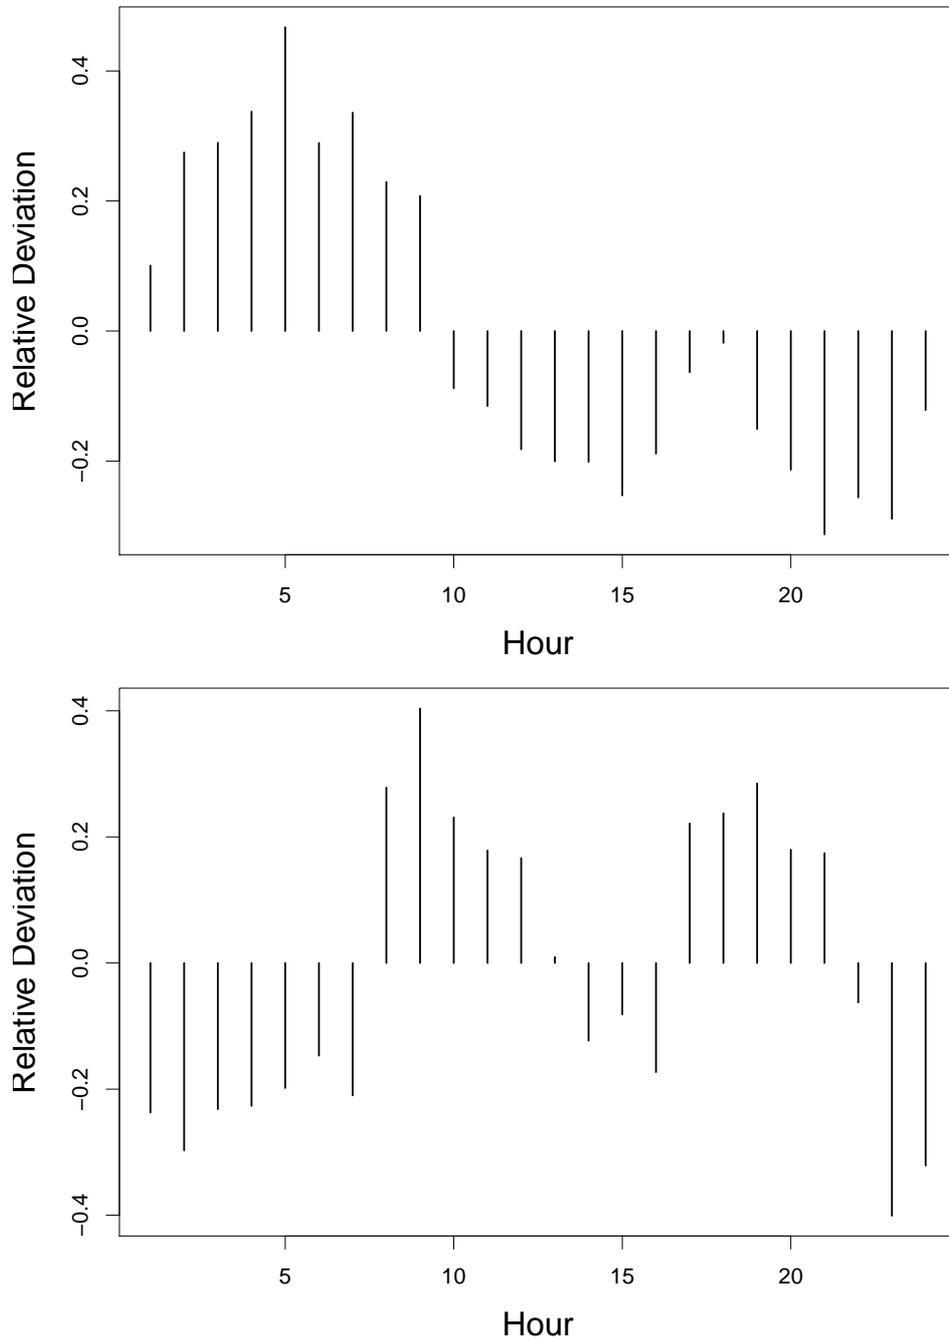The unconditional empirical probabilities of each sign for different periods are summarized in Table 1. The probabilities in the table can be seen as such that at any time the probabilities of future regulation scenarios are equal to the ones given by the table, unconditional to lead time, time in general, and other aspects of the market. Figure 14 however illustrates that there is some form for auto-correlation or persistence[1] in the sign. The plots in the figure show the percentage of hours where the imbalance sign is the same as it was the number of hours ago indicated for on the x-axis. Plots are shown for all signs combined and each of the three signs separately. Obviously the persistence is highest 1-2 hours back/ahead and is exponentially decaying, most rapidly for the no regulation and slower for the other two. A diurnal pattern is also indicated by the higher persistence around lag 24. Interestingly the persistency seems to reduce the decaying pace before reaching the unconditional empirical probabilities. This translates into that even on a day-ahead basis (13-36 hours ahead), there is a dependency on the current imbalance sign. This becomes more evident by comparing Table 2 with Table 1. Table 2 gives the averages of the 24 rightmost points of the plots in Figure 14. Since the numbers in Table 2 are consistently higher than the corresponding ones in Table 1 one must conclude that the probability of the a particular imbalance sign

---

[1]In order to separate this dependency from the conventional auto-correlation we'll term this persistency in the following

Table 1: The empirical probabilities of each imbalance sign for the full data period, the training period and the test period.

Period	Direction [%]		
	↓	~	↑
All	41.28	27.93	30.79
Training	42.80	24.74	32.46
Test	38.21	34.34	27.45

Table 2: Persistency in the imbalance sign for the full data period, the training period and the test period.

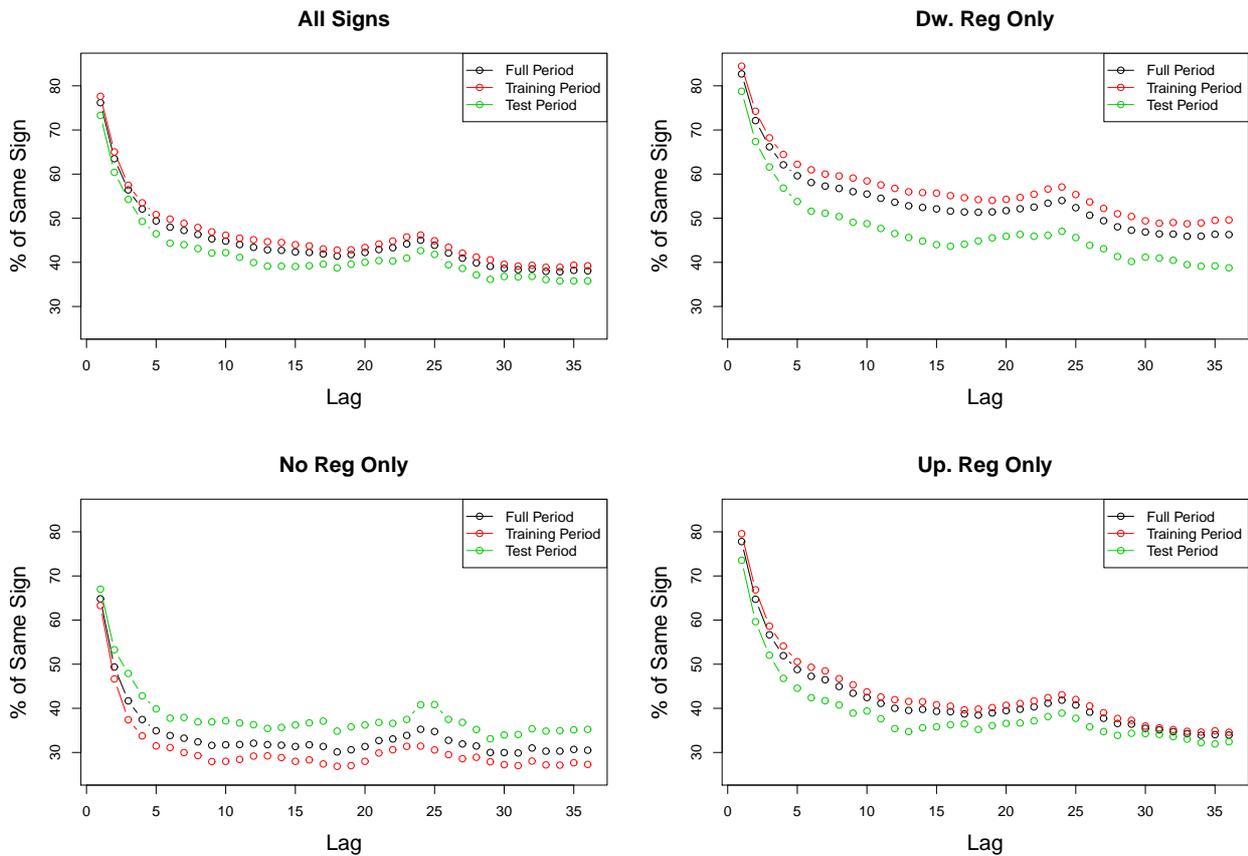| Period	Sign			
	All	↓	~	↑
All	41.10	49.96	31.60	37.86
Training	42.35	52.93	28.50	38.96
Test	38.57	43.21	36.11	35.24

Figure 14: Time persistence of the imbalance sign.

occurring again next day is higher than the overall probability of that sign.

Still working in the time-invariant domain, we now turn our attention to the impact of wind power and load on the imbalance sign. Figure 15 shows the empirical probabilities of each imbalance sign as a function of the forecast hourly wind power penetration. The plot is constructed by dividing the imbalance sign series into 10 equally populated bins according to the wind power penetration and estimating the empirical probabilities of each sign within each bin. Judging from the plot, there seems to be substantially elevated probabilities of down regulation during hours of high wind power penetration. On the contrary the probabilities of up regulation is highest during the hours with the absolutely smallest forecast wind power penetration. Also interesting to see is that the probabilities of imbalances going un-penalized decrease with increased wind power penetration.

The results presented so far on the imbalance sign have all been derived implicitly assuming time-invariant sign probabilities. However the probabilities are in fact time-varying. The first evidence of this is of course the large difference in empirical probabilities between the test period and the training period but the actual time-variance of the probabilities is probably best demonstrated by the example in Figure 16. The top two plots show an exponentially
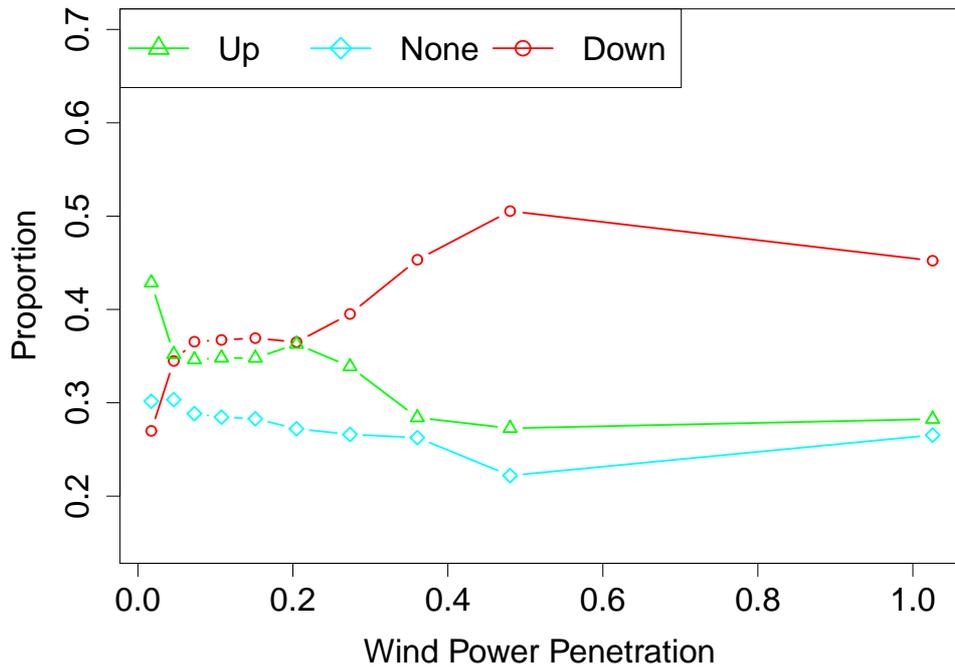
Figure 15: Empirical probabilities of $\mathbf{I}^{(\uparrow/\downarrow)}$ as a function of forecast wind power penetration

smoothed average of the actual imbalance sign series with a forgetting factor of 0.99 (left) and 0.999 (right). The bottom two plots however show an exponentially smoothed average, using the same forgetting factors, of a series that is simulated from the empirical probabilities of the full data period. Should the imbalance sign probabilities be (close to) constant in time, the smoothed averages in the top two plots would behave similarly to the averages in the bottom plots. Instead the observed averages drift much further away from the empirical probabilities for much longer time, which clearly indicates some form for non-stationary of the process.

Figure 16: Exponentially smoothed state probabilities of the observed $\boldsymbol{I}^{(\uparrow/\downarrow)}$ (top row) and the simulated $\boldsymbol{I}^{(\uparrow/\downarrow)}$ (bottom row) with $\lambda = 0.99$ (left column) and $\lambda = 0.999$ (right column). The thin lines represent the nominal probabilities for the whole period.
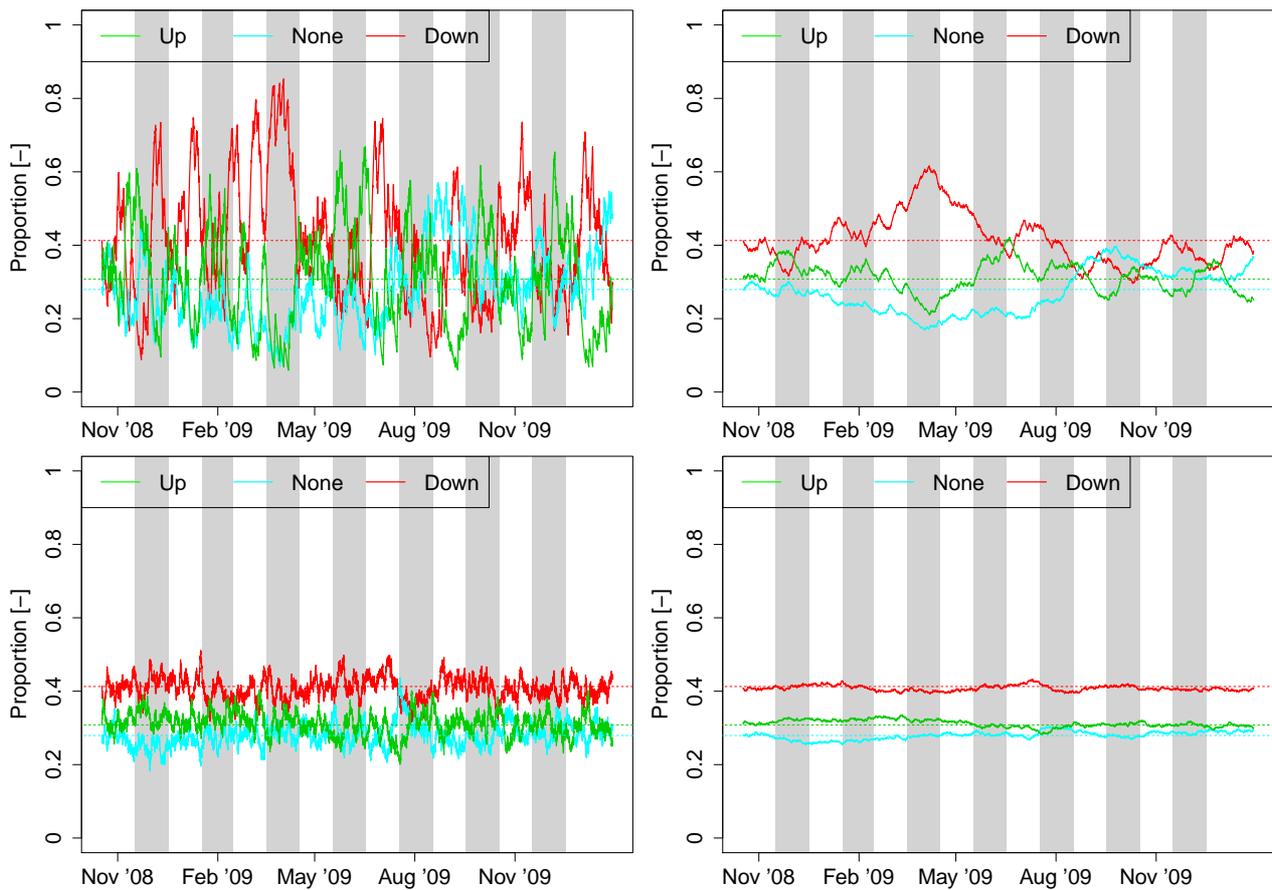
4 Holt-Winters and Conditional Holt-Winters models

Forecasts of the imbalance penalties and sign at the day-ahead level are based on the Holt-Winters model Winters [12]. This model exists in both an additive and a multiplicative version and essentially operates by updating (i) a mean component, (ii) a diurnal component, and (iii) a weekly component[2]. Each of these updates are performed at each time step and are controlled by smoothing constants α_μ, α_D, and α_W, respectively. These are all scalars between 0 and 1.

Following Gelper et al. [5], the method can be robustified against large errors by limiting the influence of the forecast error ε on the updates using the Huber influence function

$$g(\varepsilon, \tau) = \begin{cases} \varepsilon & \text{if } |\varepsilon| \leq \tau \\ \text{sign}(\varepsilon) \cdot \tau & \text{if } |\varepsilon| > \tau \end{cases} \tag{9}$$

In order to account for the effects of external variables, e.g. forecast wind power production and load, the previously described Holt-Winters model can be extended by conditioning on such variables. Let X_t be a variable on which Y_t arbitrarily depends on. Now inspired by Cleveland and Devlin [1], let x_i, $i \in [1, \ldots, M]$ be a particular *fitting point* in a set of M distinct fitting points covering the span of $\{X_t\}$. Then for each of these fitting points a Holt-Winters model is applied locally in the sense that the updates of the mean, diurnal, and weekly components for each fitting point x_i depends on the distance between x_i and X_t such that the influence of the forecast error decrease as the distance increase. This is accomplished using weights which decrease as the distance increase. The weights are assigned as

$$v_{x_i}(X_t) = V\left(\frac{||X_t - x_i||}{h(x_i)}\right) \tag{10}$$

where $V(\cdot)$ is a function taking non-negative arguments, $||\cdot||$ denotes the Euclidean norm and $h(x_i)$ is the bandwidth applied in the fitting point x_i. Following [1] and [11] a tri-cube kernel is used to determine the weights. That is

$$V(x) = \begin{cases} (1 - x^3)^3 & \text{if } x \in [0;1) \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

which entails weights between 0 and 1.

For a value of X_t with $x_i < X_t < x_{i+1}$, the forecast $\widehat{Y}_{t+k|t}$ is found by linear interpolating between $\widehat{Y}_{t+k|t,i}$ and $\widehat{Y}_{t+k|t,i+1}$. For multidimensional $\mathbf{X}_t$ and corresponding fitting points bilinear interpolation is used.

---

[2]The trend component is not used here.

The forecasting procedures outlined depend on the selection of a number of smoothing constants and bandwidths. In the following empirical study, a nearest neighbor bandwidth is applied and the corresponding smoothing constant is the fraction of observations within the bandwidth.

For imbalance penalties and the regular Holt-Winters model given N observations of $\psi^{(\downarrow)}$ and $\psi^{(\uparrow)}$, an estimate of the smoothing constants, $\widehat{\alpha} = [\ \widehat{\alpha}_\mu \ \ \widehat{\alpha}_D \ \ \widehat{\alpha}_W\ ]$ is found by minimizing the sum of squared residuals after truncating negative forecasts

$$\widehat{\alpha} = \arg\max_{\alpha} \sum_{i=1} N \left(\psi_i^{(\uparrow/\downarrow)} - \max\{0, \widehat{\psi}(\uparrow/\downarrow)_i\} \right)^2 . \tag{12}$$

For the conditional version the nearest neighbor bandwidth is included in the minimization.

For imbalance penalties and the regular Holt-Winters model given N observations of $I^{(\uparrow/\downarrow)}$, an estimate of the smoothing constants, $\widehat{\alpha} = [\ \widehat{\alpha}_\mu \ \ \widehat{\alpha}_D \ \ \widehat{\alpha}_W\ ]$, can be derived by *Maximum Likelihood*(ML). See [8, paper H] for details.

5   Modelling Results

For both the imbalance sign and the imbalance penalties model parameters are estimated for forecasting on a day-ahead basis. The day-ahead forecasts are made in such way that every day at 10:00 in the morning, forecasts are issued for the following day. Thereafter no forecasts are issued until 10:00 the following day. The day-ahead forecasts thus represent the information available before Elspot gate-closure.

5.1   Day-ahead Forecasts of Imbalance Sign

The imbalance sign forecasts are evaluated in terms of the Discrete Ranked Probability Skill Score ($RPSS_D$), calculated as

$$\begin{aligned}
RPSS_D &= 1 - \frac{\frac{1}{N} \sum_{t=1}^{T_N} RPS_{\widehat{I}_t^{(\uparrow/\downarrow)}}}{\frac{1}{N} \sum_{t=1}^{T_N} RPS_{\overline{I}_t^{\uparrow/\downarrow}}} \\
&= 1 - \frac{\sum_{t=1}^{T_N} \sum_{k=1}^{3} \left(\sum_{i=1}^{k} \widehat{I}_{t,i}^{\uparrow/\downarrow} - \sum_{i=1}^{k} I_{t,i}^{(\uparrow/\downarrow)} \right)^2}{\sum_{t=1}^{T_N} \sum_{k=1}^{3} \left(\sum_{i=1}^{k} \overline{I}_{t,i}^{\uparrow/\downarrow} - \sum_{i=1}^{k} I_{t,i}^{(\uparrow/\downarrow)} \right)^2 + D}
\end{aligned} \tag{13}$$

where $RPS_{\widehat{I}_t^{(\uparrow/\downarrow)}}$ and $RPS_{\overline{I}_t^{\uparrow/\downarrow}}$ are the *Ranked Probability Scores* [4, 9, 10] of the estimated posterior probabilities and the climatology forecasts respectively. Furthermore, D is a

Table 3: Estimated parameters and $RPSS_D$ for the different model setups investigated

Model		$h(v)$	α_μ	α_{24}	α_{168}	$RPSS_D$ Training set	Test set
Unconditional	Mdl 1	—	0.0020	—	—	0.5117	0.5019
	Mdl 2	—	0.0130	0.1067	—	0.5206	0.5054
	Mdl 3	—	0.0197	—	0.0689	0.5165	0.5007
	Mdl 4	—	0.0146	0.1104	0.0477	0.5207	0.5035
Conditional	Mdl 1	0.9505	0.0248	—	—	0.5045	0.4885
	Mdl 2	0.9858	0.0124	0.4785	—	0.5193	0.5049
	Mdl 3	0.8621	0.0558	—	0.1587	0.5163	0.5012
	Mdl 4	0.9543	0.0156	0.4898	0.0829	0.5178	0.5024
Conditional 2D	Mdl 1	0.9305	0.0274	—	—	0.5004	0.4890
	Mdl 2	0.9666	0.0140	0.4877	—	0.5188	0.5042
	Mdl 3	0.9816	0.0416	—	0.1378	0.5161	0.5026
	Mdl 4	0.9782	0.0104	0.4780	0.0728	0.5177	0.5036

bias correction term found as

$$D = \frac{1}{N} \sum_{k=1}^{3} \sum_{i=1}^{k} \left[p_i \left(1 - p_i - 2 \sum_{j=i+1}^{k} p_j \right) \right] \tag{14}$$

Table 3 summarizes the estimated parameters and the corresponding $RPSS_D$ for the model setups tried. Also, the table show the result of a setup where load- and the wind power forecasts enter into the model as separate explanatory variables (Conditional 2D).

The reliability of the sign forecasts can be estimated by segmenting the forecasted probabilities into a small number of equally populated bins and taking the mean of each bin as the predicted nominal probability. This average predicted probability is then compared to the empirical frequencies in each group. This is further discussed in [8, paper H], where a similar traning period, but a longer test period is used. Generally, the predicted probabilities seems reliable.

5.2 Day-ahead Forecasts of Imbalance Penalties

The imbalance penalty forecasts are evaluated in terms of the Root Mean Square Error (RMSE). For the sake of comparison, Table 4 show the standard deviation of the raw series. These values can be intrepreted as the RMSE assuming the forecast to be the overall mean of the respective series.

For the day-ahead forecasts six different conditioning setups are tried and as before, four

Table 4: Standard Deviation of the down-and up-regulation penalties

Period	Down	Up
Training	84.5987	77.5868
Test	170.4963	119.6455

different combinations of models terms both in the additive and the multiplicative formulation. These conditioning setups are:

- Unconditional HW

- Conditional to the forecast Spot Price

- Conditional to the forecast Load

- Conditional to the forecast Wind Power Penetration

- Conditional to the forecast Spot Price and Load

- Conditional to the forecast Spot Price and Wind Power Penetration

In order to save space only the results for the few best performing models are reported here. Tables 5 & 6 list the estimated model parameters and the corresponding *RMSE* for the up-and down-regulation penalty models respectively. The best performing models are conditional ones, all of which involve conditioning upon the forecast spot price. Moreover, the models include either solely a mean term or a mean and a daily seasonal term. None of the models however perform particularly well and their residual *RMSE* is only marginally less than the series standard deviation of the series. In light of the optimal smoothing parameters, which yield long model memory (small values of the smoothing constants), this is not surprising. However, in context with the presumed non-stationarity of the penalties this long memory was not expected.

Figures 17 & 18 show the residual histograms for the up-regulation penalty forecasts during the training and the test period respectively. On the left the residuals for the forecasts made with the HW-model conditioned on the spot price and the load forecast with multiplicative daily seasonality are shown. To the right, a histogram of the penalties with the mean subtracted is given. The same plots for the down-regulation penalties are given in Figures 19 & 20.

Figures 21 & 22 show the cummulative density function (CDF) for the forecast residuals and the penalty series to the left and sharpness of the centred quantile ranges of the same to the right. The bottom two plots of the each figure show the same only for the absolute residuals.

Table 5: Model parameters and RMSE for the best performing up-regulation penalty models

Seasonality	X	$h(\cdot)$	τ	α_μ	α_D	α_W	RMSE Train	Test
	$\pi^{(S)}$ & L	0.3301	5469	0.0059	–	–	72.7742	113.3243
	$\pi^{(S)}$	0.1136	358079	0.0184	–	–	73.0819	113.6558
Additive	$\pi^{(S)}$ & L	0.4998	31989	0.0031	0.0347	–	73.3681	113.1332
	$\pi^{(S)}$	0.4317	57316	0.0014	0.0649	–	74.1275	114.0328
	$\pi^{(S)}$ & W	0.4500	59589	0.0007	0.0622	–	73.8089	114.5221
Multiplicative	$\pi^{(S)}$ & W	0.0957	1354	0.0125	0.0707	–	73.8830	113.1240
	$\pi^{(S)}$ & L	0.4440	3762	0.0015	0.0349	–	73.2870	113.1567
	$\pi^{(S)}$	0.3922	1538	0.0018	0.0437	–	73.6592	113.8121

Table 6: Model parameters and RMSE for the best performing down-regulation penalty models

Seasonality	X	$h(\cdot)$	τ	α_μ	α_D	α_W	RMSE Train	Test
	$\pi^{(S)}$ & L	0.0985	848	0.1174	–	–	69.2965	167.4267
	$\pi^{(S)}$ & W	0.1093	6448	0.0442	–	–	70.5603	168.7917
Additive	$\pi^{(S)}$ & W	0.8139	1994	1.45×10^{-6}	0.1222	–	70.5076	168.3655
	$\pi^{(S)}$ & L	0.7299	2480	0.0012	0.0767	–	71.8044	169.1354
Multiplicative	$\pi^{(S)}$ & W	0.1189	∞	0.0542	0.0644	–	69.7495	165.6681
	$\pi^{(S)}$ & L	0.1013	384	0.0979	0.0803	–	68.8445	169.8900

Combined these six figures illustrate that the forecasting skill the conditional HW-model have is to a certain degree owed to that they reduce a bit the largest residuals. That in it self is a good thing although a better general forecasting ability would be appreciated.

Finally, Figures 23 & 24 show the empirical autocorrelation functions (ACFs) for the residuals and the squared residuals for the up- and down-regulation penalties respectively. Obviously, there is considerable auto-correlation left in the residuals, especially for the down regulation penalties. However, most of this auto-correlation lies on the intra-day horizons so maybe there are limited possibilities to utilize it for day-ahead forecasts.
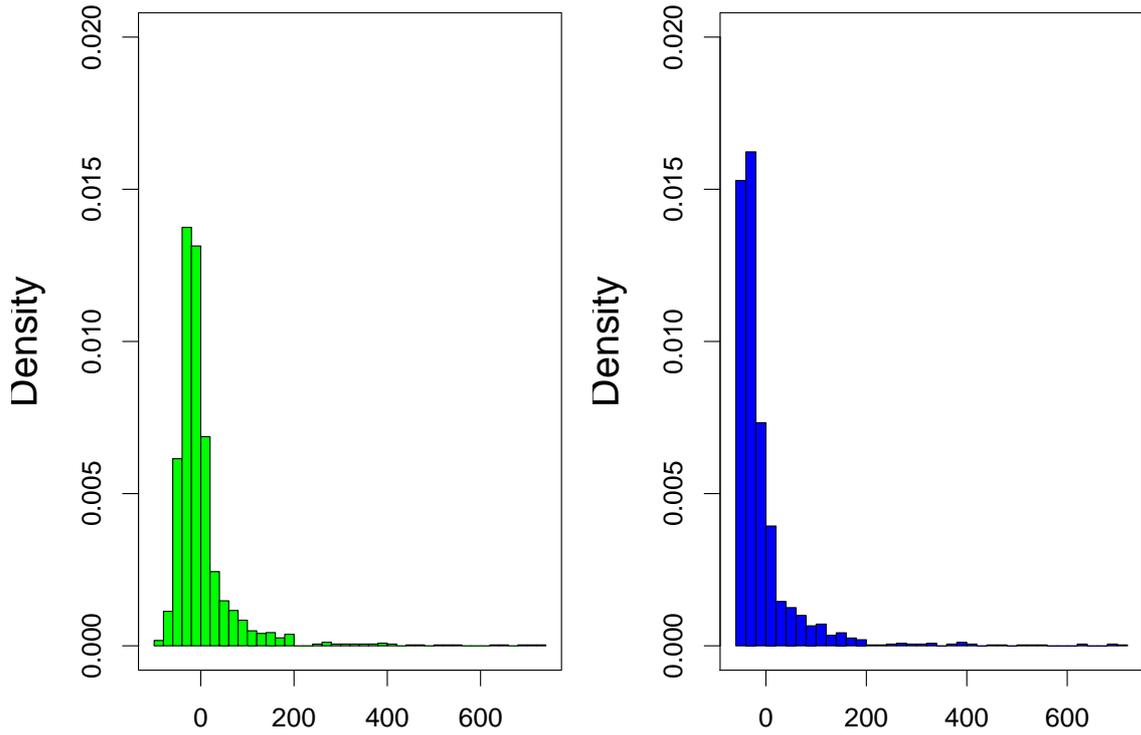
Figure 17: Residual histograms for the up-regulation penalties during the training period. The residuals of the HW-model are represented on the left while the residuals of the constant mean prediction are represented on the right
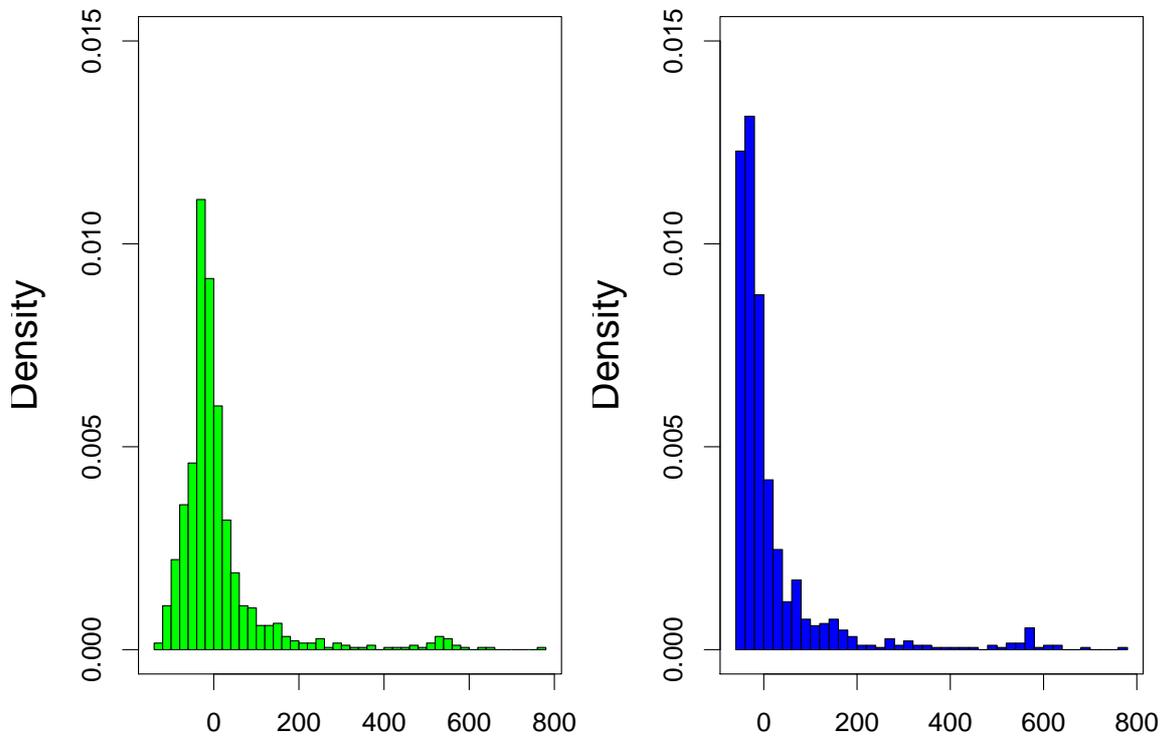

Figure 18: Residual histograms for the up-regulation penalties during the test period. The residuals of the HW-model are represented on the left while the residuals of the constant mean prediction are represented on the right
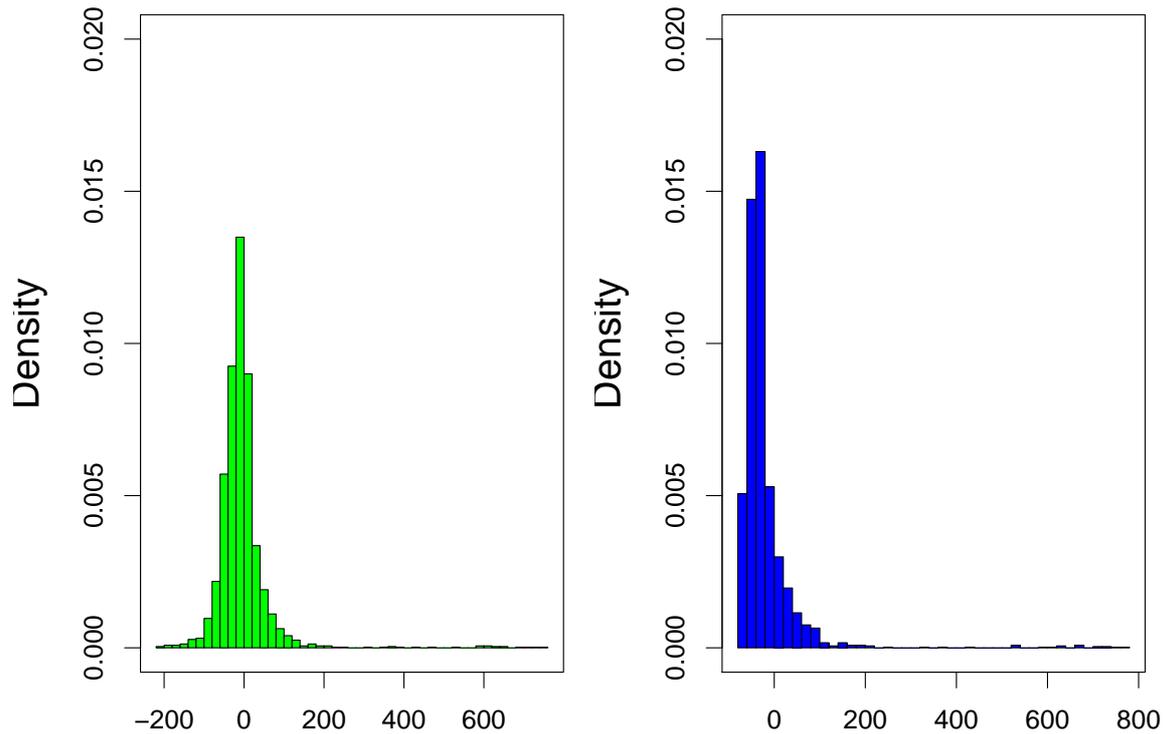
Figure 19: Residual histograms for the down-regulation penalties during the training period. The residuals of the HW-model are represented on the left while the residuals of the constant mean prediction are represented on the right
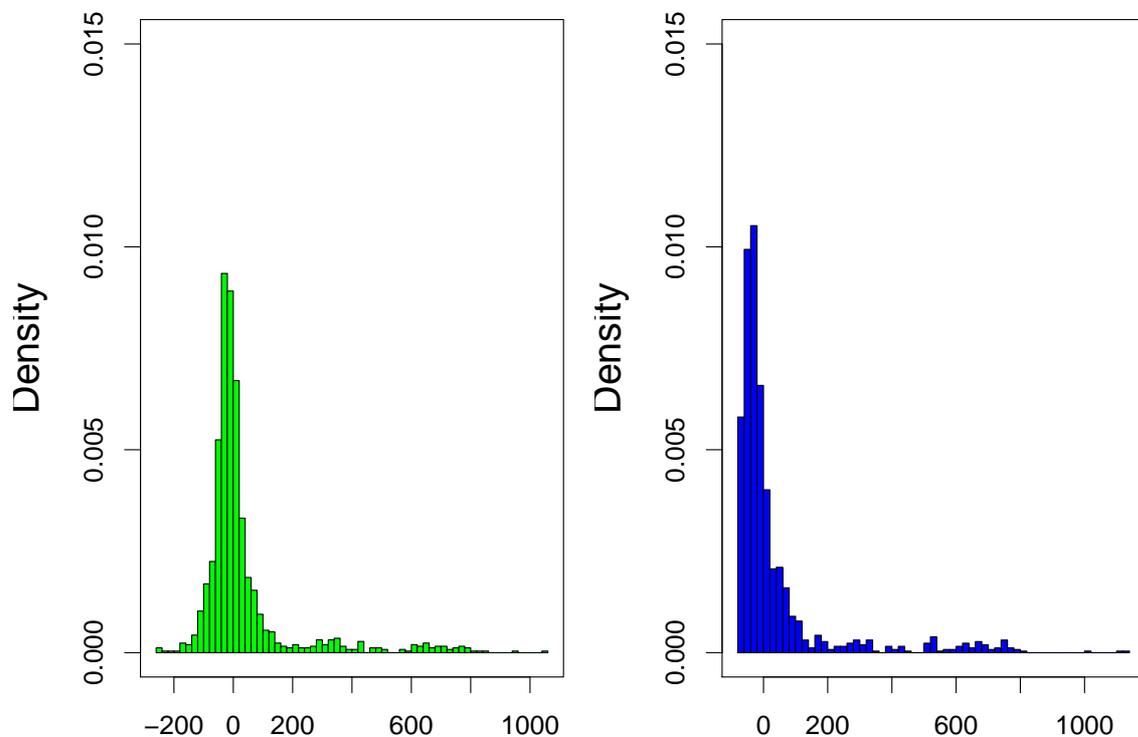

Figure 20: Residual histograms for the down-regulation penalties during the test period. The residuals of the HW-model are represented on the left while the residuals of the constant mean prediction are represented on the right
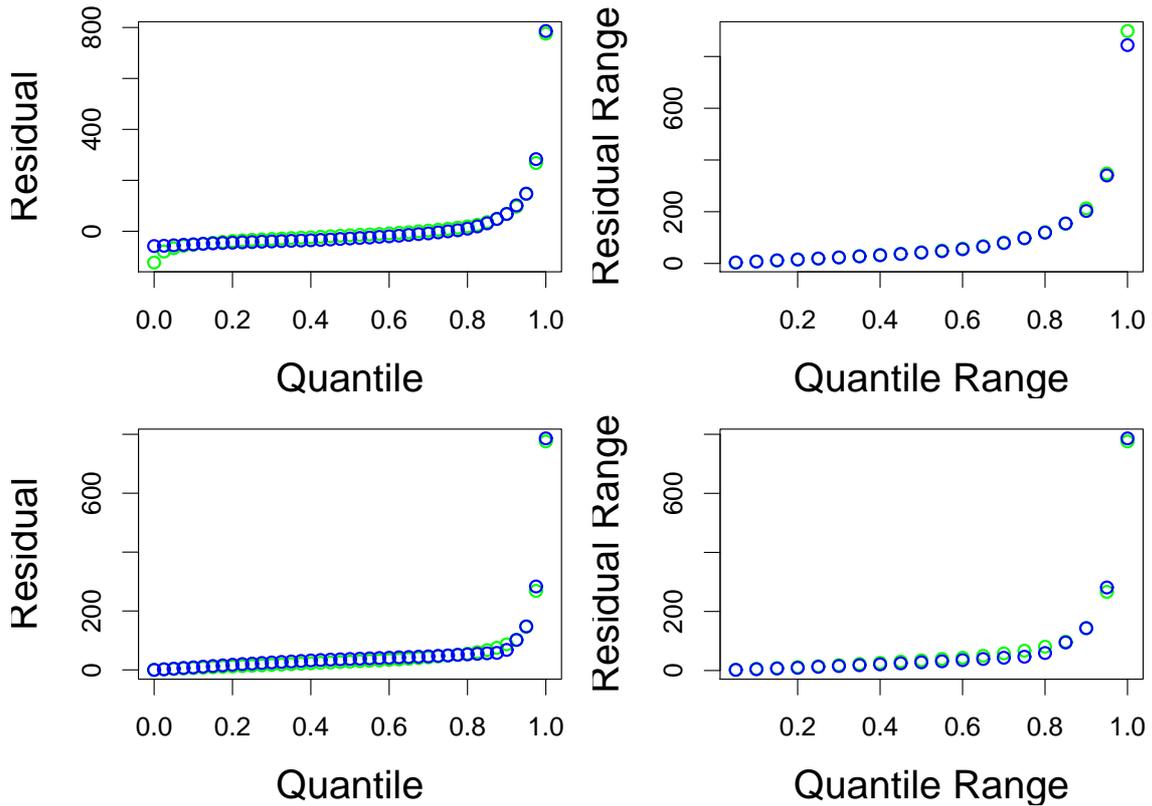
Figure 21: The CDF and the centered sharpens of the residuals (top) and the absolute residuals (bottom) for the up regulation penalties
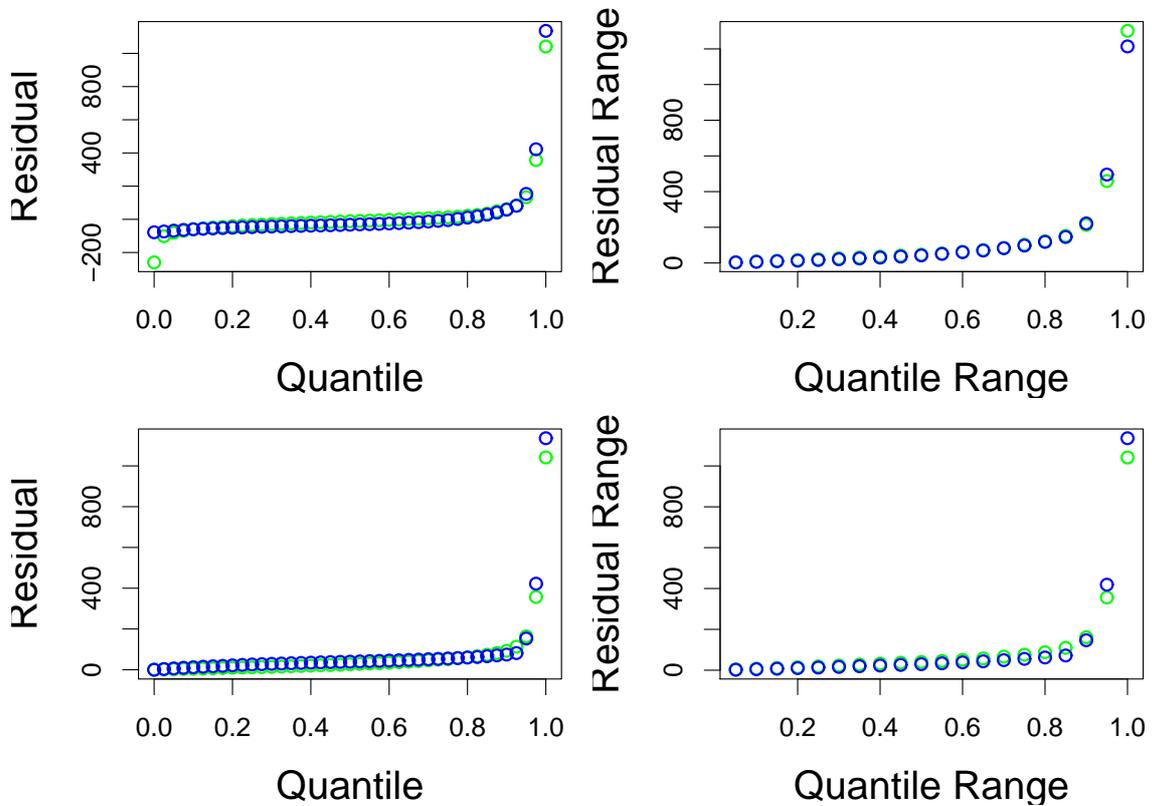

Figure 22: The CDF and the centered sharpens of the residuals (top) and the absolute residuals (bottom) for the down regulation penalties
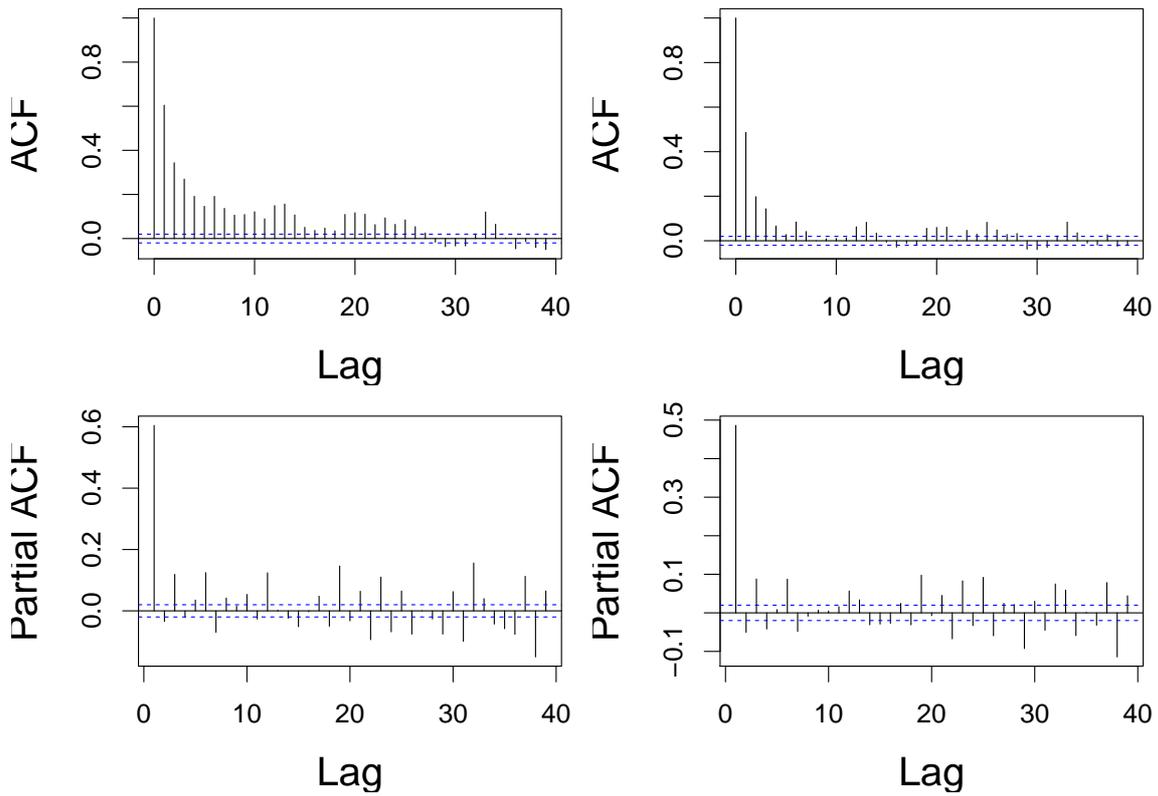
Figure 23: ACFs for the residuals (left) and the squared residuals (right) for the up regulation penalties
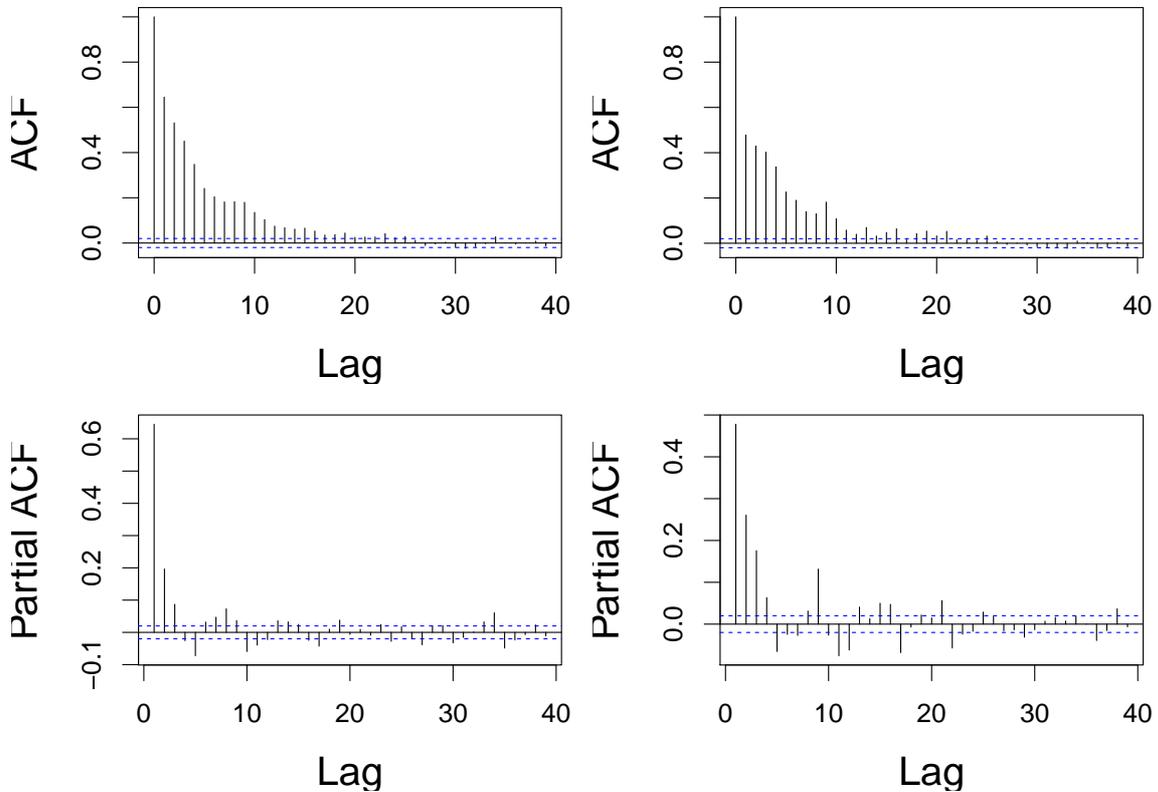

Figure 24: ACFs for the residuals (left) and the squared residuals (right) for the down regulation penalties

5.3   Day-ahead forecasts of unconditional expectations

The forecasts of imbalance sign and imbalance penalties considered above in principle contain the full information available at forecast time. However, as mentioned in Section 2 the unconditional expectations of the up and down regulation penalties are needed in certain bidding strategies. Furthermore as described in the analysis of the FlexPower setup ENFOR/08EKS0004A001-A[2] only the expected price is required.

The relation between imbalance sign probabilities, conditional- and unconditional-expectations of the imbalance penalties are stated in the end of Section 2. The unconditional expectation of the final price for the hour π_t can be expressed as

$$\mathbb{E}[\pi_t] = \mathbb{E}[\pi_t^{(S)} + \psi_t^{(\uparrow)} - \psi_t^{(\downarrow)}] = \mathbb{E}[\pi_t^{(S)}] + \mathbb{E}[\psi_t^{(\uparrow)}] - \mathbb{E}[\psi_t^{(\downarrow)}].$$

Where as explained in Section 2 at least one of $\psi_t^{(\uparrow)}$ and $\psi_t^{(\downarrow)}$ will be zero, while this is not true for the expected values of these. For horizons up to approximately 12 hours the spot price $\pi_t^{(S)}$ will be known and hence $\mathbb{E}[\pi_t^{(S)}] = \pi_t^{(S)}$.

Since the unconditional expectations of the imbalance penalties has an important role in relation to many decision problems, including those related to FlexPower, special attention is put on evaluation of these. Based on the day-ahead forecasts of imbalance penalties and imbalance signs the unconditional expectations of the imbalance penalties can be calculated as described above. This result in a forecast of $\psi_t^{(\uparrow)}$ and $\psi_t^{(\downarrow)}$ for each hour of a day available at the day-ahead level, i.e. in the morning before the spot prices has been fixed.

These forecast could be evaluated directly against the actual imbalance penalties (including the zeros). However, as has previously been established the noise level is large and therefore we will rather attempt to evaluate the expected values. This is accomplished by grouping the forecast unconditional imbalance penalties into 10 groups and calculating the mean of the actual imbalance penalties within each group. Plotting the actual means against the means of the forecasts will reveal if the unconditional means are appropriately forecasted. The result is displayed in Figure 25.

5.4   Intra day forecasts of unconditional expectations

In order to simplify the models applied for intra-day forecasting of imbalance penalties models directly modelling the unconditional expectations are developed. One disadvantage of this approach is that it does not fully model the imbalance market. However, if unconditional means are only required in the decision/optimization process the modelling process is simpler.
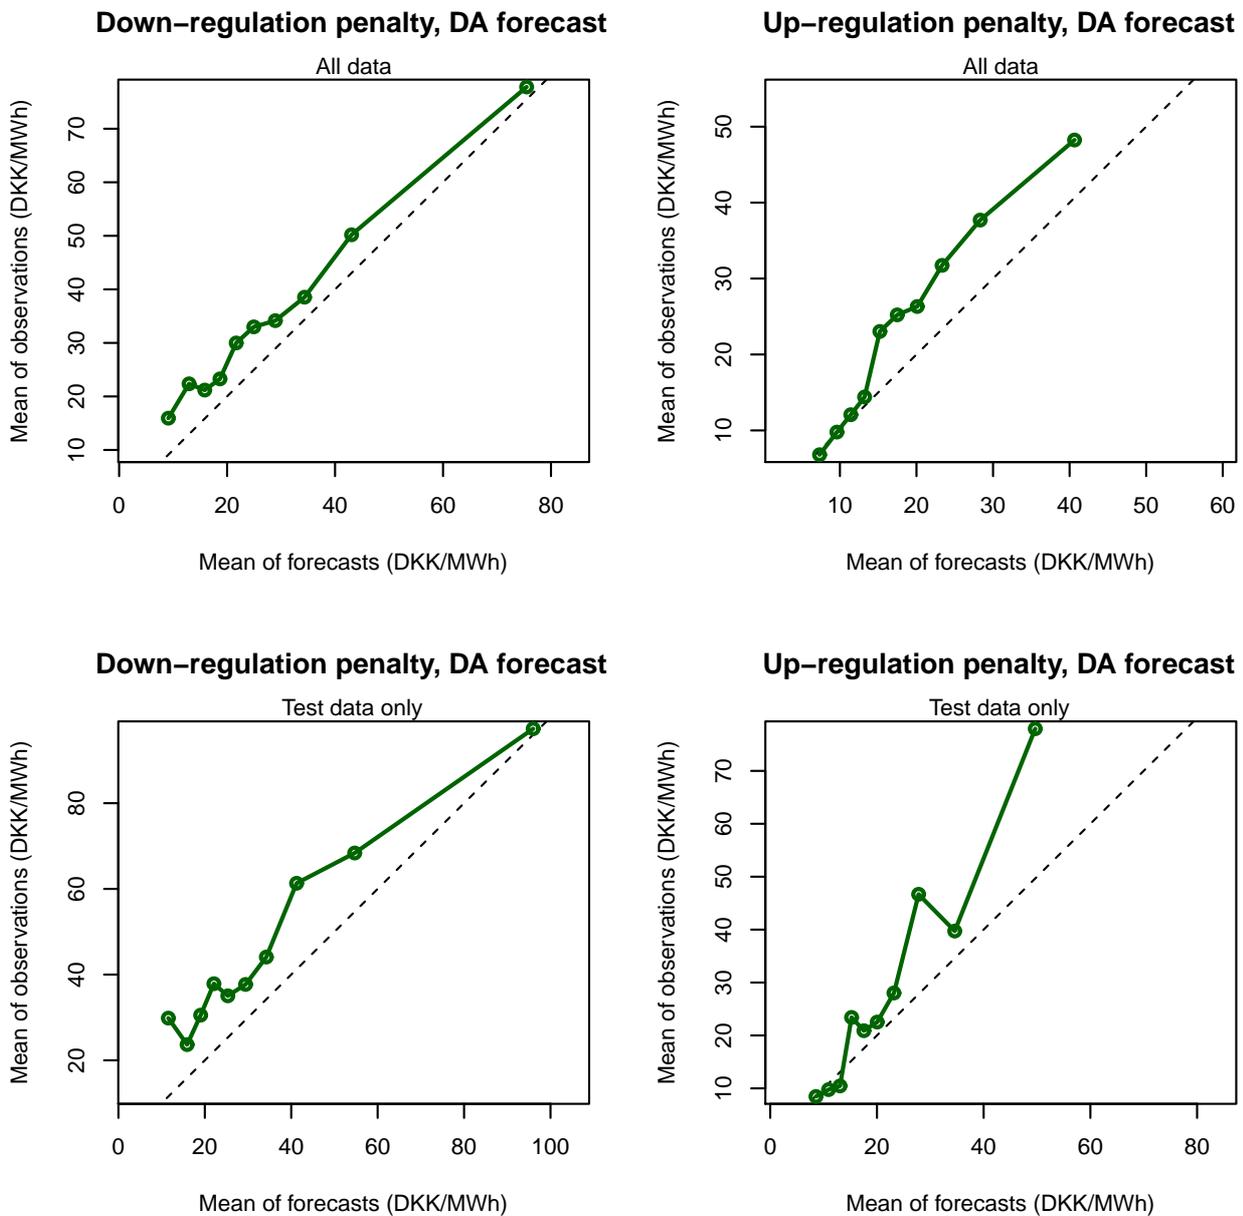
Figure 25: Evaluation of day-ahead forecasts of unconditional imbalance penalties.

Focus is on horizons up to 12 hours where the spot price is always known as mentioned above. The models are auto-regressive models applied to each horizon separately. The models includes a diurnal term, i.e. the actual imbalance penalty 24 hours before the time point for which a forecast is required.

The models are estimated adaptively. This allows the estimated models to track changes in the dynamics of the imbalance market. The adaptivity is controlled by the forgetting factor
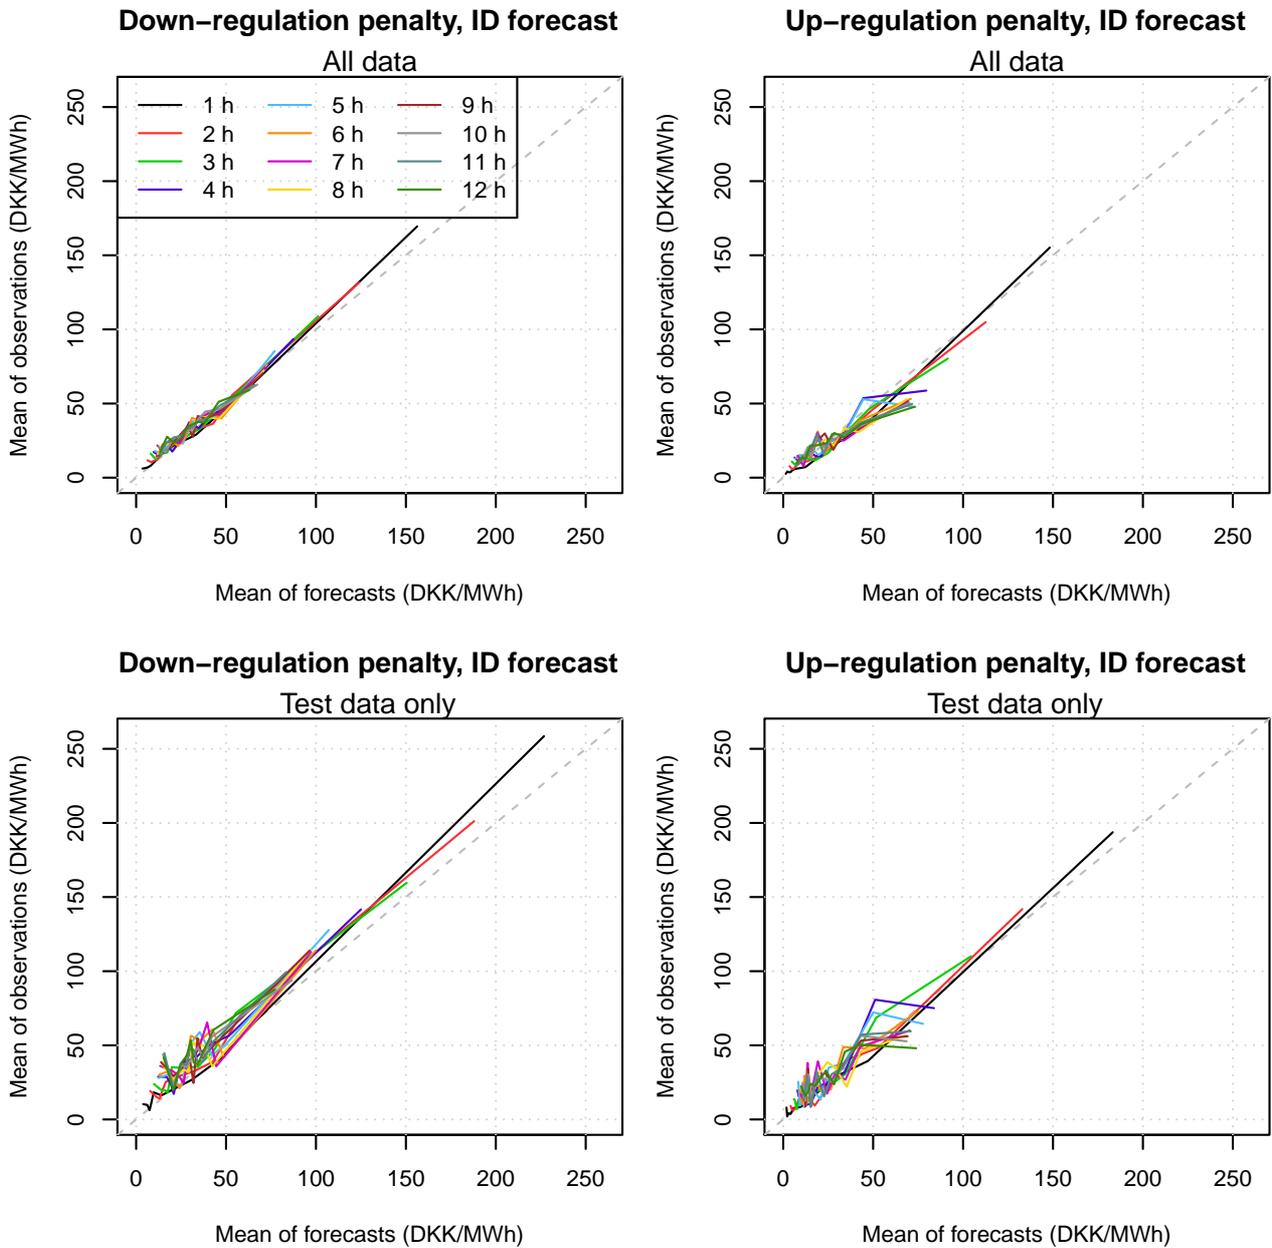
Figure 26: Evaluation of intra-day forecasts of unconditional imbalance penalties. Note that the intervals covered by the axes are different from those in Figure 25.

for which an exhaustive search has not been performed. However, a forgetting factor of approximately six weeks seems appropriate. The adaptive estimation were initiated based on data already from 2008-01-01, but evaluation periods for both "all data" and "test data" are as stated previously. This means that for a forgetting factor corresponding to six weeks all effects of initialization has vanished at the beginning of the evaluation periods used. Figure 26 depicts diagnostic plots for forecast horizons 1-12 hours.

6 Conclusion

This report considers forecasting of imbalance penalties for the Western Danish price area of Nord Pool. Specifically, the period November 2008 to January 2010 (15 months) is considered. Two-thirds of the period (November 1st 2008 – September 10th 2009) is used as training period and one-third (September 11th 2009 – January 31st 2010) is used as testing period. A number of explanatory variables are considered and the report mainly focus on the day-ahead horizons.

The report consider what is here called down- and up-regulation imbalance penalties, i.e. the positive differences between the spot and the down- and up-regulation prices, respectively. Since, at any particular point in time, at most one regulation price can be different from the spot price the imbalance penalties will often be exactly zero. Therefore the report considers separately the tasks of forecasting the probability of a particular imbalance sign (down, up, or no regulation penalty) and the magnitude of the regulation penalty given that it is strictly positive. The last value is called the conditional expectation, i.e. the expected value conditional on the penalty being positive. However, for decision making based on expected revenues, the unconditional expectations of the down- and up-regulation penalties are required. The report describes how these unconditional expectations can be found from the imbalance sign probabilities and the conditional expectations.

Hardly surprising, plotting the data reveals that the imbalance penalties, especially the up-regulation imbalance penalty, contain large spikes. Nevertheless, the imbalance penalties seem to contain a diurnal variation both with respect to the sign and the magnitude. Furthermore, the sign-probabilities show some dependence on the wind power forecast for the region.

For the imbalance sign probabilities a number of models both with and without explanatory variables are considered. However, the benefit from using explanatory variables is not clear. For the magnitude of the penalties there seems to be some benefit from conditioning on the forecast spot price. However, the fundamental noise level at day-ahead is dominating.

As mentioned above the unconditional expectations of the imbalance penalties are important for decision making based on expected values. Consequently these forecasts are evaluated also. The evaluation is performed by splitting the forecast values in ten equally populated bins and calculating the means in each group. Representing each group by the mean of the forecasts belonging to this group and plotting these against the observed means results in ten points which, except for random variation, should lie on the line of identity. Furthermore, an AR-model is used directly to forecast the unconditional expected imbalance penalties for horizons up to 12 hours and the results are evaluated using the same principle as just described. In both cases the observed and forecasted means lie close to the line of identity. See ENFOR/08EKS0004A003-A[3] for further analyses of this kind, including the an evaluation of the uncertainty of the point location. Note that in that reference the penalties are called imbalance unit costs.

References

[1] William S. Cleveland and Susan J. Devlin. Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association*, 83(403): 596–610, 1988.

[2] ENFOR/08EKS0004A001-A. Forecast requirements for house temperature control with flexible energy prices, 2011.

[3] ENFOR/08EKS0004A003-A. Simulation of 5 minute prices based on actual 1 hour data, 2013.

[4] E Epstein. A scoring system for probability forecasts of ranked categories. *Journal of Applied. Meteorology*, 8:985–987, 1969.

[5] Sarah Gelper, Roland Fried, and Cristophe Croux. Robust forecasting with exponential and Holt-Winters smoothing. *Journal of Forecasting*, 29(3):285–300, 2010.

[6] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer, 2001.

[7] Tryggvi Jónsson. Mean-CVaR bidding strategy for both price taking and price making wind power producer. Technical report, 2011.

[8] Tryggvi Jónsson. *Forecasting and decision-making in electricity markets with focus on wind energy*. PhD thesis, Technical University of Denmark, 2012. IMM-PHD-2012.

[9] Allan H. Murphy. On the ranked probability skill score. *Journal of Applied Meteorology*, 8:988–989, 1969.

[10] Allan H. Murphy. A note on the ranked probability skill score. *Journal of Applied Meteorology*, 10:155–156, 1971.

[11] Henrik Aa. Nielsen, Torben S. Nielsen, Alfred K. Joensen, Henrik Madsen, and Jan Holst. Tracking time-varying coefficient functions. *International Journal of Adaptive Control and Signal Processing*, 14(8):813–828, 2000.

[12] Peter R. Winters. Forecasting sales by exponentially weighted moving averages. *Management Science*, 6(3):324–342, 1960.

[13] Marco Zugno, Pierre Pinson, and Tryggvi Jónsson. Trading wind energy based on probabilistic forecasts of wind generation and market quantities. *Wind Energy*, submitted, 2011.